

On the Complexity of Computing with Planar Algebraic Curves

Alexander Kobel¹⁻³

Michael Sagraloff¹

¹Max-Planck-Institut für Informatik

²International Max Planck Research School for Computer Science

³Universität des Saarlandes

Saarbrücken, Germany

{alexander.kobel,michael.sagraloff}@mpi-inf.mpg.de

1st August, 2014

Abstract

In this paper, we give improved bounds for the computational complexity of computing with planar algebraic curves. More specifically, for arbitrary coprime polynomials $f, g \in \mathbb{Z}[x, y]$ and an arbitrary polynomial $h \in \mathbb{Z}[x, y]$, each of total degree less than n and with integer coefficients of absolute value less than 2^τ , we show that each of the following problems can be solved in a deterministic way with a number of bit operations bounded by $\tilde{O}(n^6 + n^5\tau)$, where we ignore polylogarithmic factors in n and τ :

- *The computation of isolating regions in \mathbb{C}^2 for all complex solutions of the system $f = g = 0$,*
- *the computation of a separating form for the solutions of $f = g = 0$,*
- *the computation of the sign of h at all real valued solutions of $f = g = 0$, and*
- *the computation of the topology of the planar algebraic curve \mathcal{C} defined as the real valued vanishing set of the polynomial f .*

Our bound improves upon the best currently known bounds for the first three problems by a factor of n^2 or more and closes the gap to the state-of-the-art randomized complexity for the last problem.

Keywords. polynomial system solving, bivariate systems, separating form, topology analysis, algebraic curves, arrangement computation, cylindrical algebraic decomposition, complexity analysis, multipoint evaluation, approximate computation

1 Introduction

In this paper, we derive record bounds for the computational complexity of the following problems, which are related to the arrangement computation of planar algebraic curves:

- (P1) Given two coprime polynomials $f, g \in \mathbb{Z}[x, y]$ of degree n or less, compute isolating regions in \mathbb{C}^2 for *all distinct complex* solutions $(x_i, y_i) \in \mathbb{C}^2$ of the system

$$f(x, y) = g(x, y) = 0, \quad (1)$$

with $i = 1, \dots, r$ and some integer r with $r \leq \deg f \cdot \deg g \leq n^2$, which is the upper bound of the number of solutions of the zero-dimensional system due to Bézout's Theorem.

- (P2) Compute a *separating form* $x + s \cdot y$, with $s \in \{0, 1, \dots, n^4\}$, for (1) such that $x_i + s \cdot y_i \neq x_j + s \cdot y_j$ for all i, j with $i \neq j$.

- (P3) Given an arbitrary polynomial $h \in \mathbb{Z}[x, y]$, evaluate the sign of h at all *real valued* solutions (x_i, y_i) of (1).

- (P4) Given an arbitrary polynomial $f \in \mathbb{Z}[x, y]$, compute the topology of the real planar algebraic curve

$$\mathcal{C} := \{(x, y) \in \mathbb{R}^2 : f(x, y) = 0\} \quad (2)$$

in terms of a planar straight line graph that is isotopic¹ to \mathcal{C} .

We remark that a solution to the above problems allows us to answer all necessary queries for arrangement computations with planar algebraic curves [7, 8, 21, 25, 31]. Namely, for a set of planar algebraic curves, we can compute the topology of each of these curves from (P4), we can compute the intersection points of two curves from (P1), and, from (P3), we can decide whether two intersection points from two distinct pairs of curves are equal or not.

The main contribution of this paper is a *deterministic* algorithm that solves all of the above problems (P1) to (P4) in a number of bit operations bounded by $\tilde{O}(n^6 + n^5\tau)$, where n is an upper bound for the total degree of the polynomials f, g and h , and τ is an upper bound for the bitsize of their coefficients. For the first two problems, we also give more general bounds that take into account the case of unbalanced input. That is, if f and g are polynomials of total degree m and n and with integer coefficients of bitsize bounded by τ_f and τ_g , respectively, then (P1) and (P2) can be solved in time $\tilde{O}(\max\{m, n\} \cdot (m^2n^2 + mn(m\tau_g + n\tau_f)))$.

We briefly outline our approach: For (P1), we first extend (and modify) an algorithm from Berberich et al. [8, 6], denoted BiSOLVE, that isolates only the *real valued* solutions of (1). The so-obtained algorithm CBiSOLVE computes isolating polydisks in \mathbb{C}^2 for all complex solutions and further refines these disks to an arbitrarily small size if necessary. From a high-level perspective, the algorithm decomposes into two steps: In the first step, the *projection step*, we project all solutions onto their x - and y -coordinates using resultant computation and univariate root finding. This induces a grid consisting of $O(n^4)$ candidates that have to be checked for solutions in the second step, the *validation step*: For processing the

¹We consider the stronger notion of an *ambient isotopy*, but omit the word “ambient”. A graph \mathcal{G}_C , embedded in \mathbb{R}^2 , is ambient isotopic to C if there exists a continuous mapping $\phi : [0, 1] \times \mathbb{R}^2 \mapsto \mathbb{R}^2$ with $\phi(0, \cdot) = \text{id}_{\mathbb{R}^2}$, $\phi(1, C) = \mathcal{G}_C$, and $\phi(t_0, \cdot) : \mathbb{R}^2 \mapsto \mathbb{R}^2$ is a homeomorphism for each $t_0 \in [0, 1]$.

candidates, we combine approximate evaluation of the input polynomials at the candidates and adaptive evaluation bounds derived from the co-factor representations of the resultant polynomials; see Section 2.3 for details. We further remark that CBISOLVE does not need any coordinate transformation and returns isolating polydisks in the initial coordinate system.

From the solutions of the system (1), we derive a corresponding separating form $x + s \cdot y$ for (1) by approximating all “bad” values for s (to an error of $1/2$), for which a pair of distinct solutions is mapped to the same value via $x + s \cdot y$. Since there are at most r , with $r \leq n^2$, many distinct solutions, there exist at most $\binom{r}{2} < n^4$ bad values for s . Hence, we can determine a separating form with an $s \in \{0, 1, \dots, n^4\}$. This solves Problem (P2).

For Problem (P3), we first isolate the solutions of the two systems $f = g = 0$ and $g = h = 0$ and, then, determine all common solutions by comparing the corresponding coordinate values. For the sign evaluation of h at the remaining solutions of $f = g = 0$, we again use approximate evaluation.

For Problem (P4), we first perform a coordinate transformation that ensures that the curve \mathcal{C} is in generic situation with respect to the critical points² of its defining polynomial f , that is, there exist no two critical points of f sharing the same x -coordinate. In particular, this guarantees that, for any $\alpha \in \mathbb{R}$, there exists at most one singular point of \mathcal{C} located above α . We remark that this approach crucially differs from previous approaches that require a coordinate transformation that is generic with respect to all x -critical points³ of the curve. In a second step, we compute all *real valued* x -critical points $(x_i, y_i) \in \mathbb{R}^2$ of \mathcal{C} , with $i = 1, \dots, k$ and some $k \in \mathbb{N}$, and evaluate the sign of $f_x := \frac{\partial f}{\partial x}$ at all points (x_i, y_i) . Based on the latter computation, we can use Teissier’s Lemma [8, 30, 51] in order to compute the number of distinct (complex) roots of each of the polynomials $f(x_i, y) \in \mathbb{R}[y]$; see Section 3 for details. Given the latter information, we can further isolate the roots of the polynomials $f(x_i, y)$ using a certified numerical method [39] that works with approximations of the polynomials $f(x_i, y)$ only. We further choose values $\gamma_i \in \mathbb{R}$ that separate the x -critical values x_i and isolate the roots of the square-free polynomials $f(\gamma_i, y)$. Eventually, we connect the “lifted points” (i.e. points $(x, y) \in \mathcal{C}$, where $x = x_i$ or $x = \gamma_i$) via line segments in an appropriate manner. We remark that this last step can be achieved purely based on combinatorial decisions as we can derive the local topology at all non-singular x -critical points, and there are no two singular points of the curve located above each other; see Section 3 for details.

The analysis of CBISOLVE is based on and improves the results by Emeliyanenko and Sagraloff from [24]. In particular, the worst case bit complexity of CBISOLVE improves by a factor of n^2 , due to the use of asymptotically fast methods for approximate multipoint evaluation [33, 35] and univariate root finding [39, 42, 46], and also due to a slight, but crucial, modification of BISOLVE in the validation step. For the analysis of our algorithms for solving a bivariate system, for computing a corresponding separating form, and for the sign evaluation of a polynomial at the real-valued solutions, we heavily utilize amortization over the set of all solutions. For instance, our analysis shows that the cost for computing the sign of a polynomial h at only one solution of (1) can, in the worst case, be of the same order as the cost for computing the sign of h at all solutions. For analyzing our algorithm for topology computation, we use amortized complexity bounds as provided in [39, 32].

²A point $(x, y) \in \mathbb{C}^2$ is critical for the function f if $\frac{\partial f}{\partial x}(x, y) = \frac{\partial f}{\partial y}(x, y) = 0$.

³A point $(x, y) \in \mathbb{C}^2$ is an x -critical point of the curve \mathcal{C} if $\frac{\partial f}{\partial x}(x, y) = f(x, y) = 0$.

In fact, we consider our algorithm as proposed in the paper at hand to be a deterministic variant of the randomized method introduced in [8] and analyzed in [39]. Here, we remark that the key idea that allows us to derandomize the algorithm from [8] is to consider a coordinate transformation that is generic with respect to the critical points of f and not (necessarily) with respect to the x -critical points of \mathcal{C} . Following this approach, we can efficiently solve three major problems with respect to topology computation, that is, the computation of a generic coordinate transformation, the computation of the number of roots of the polynomials $f(x_i, y)$, where x_i is an x -critical value, and the computation of the adjacencies between neighboring fibers in the final step of our algorithm.

Related work Computing the solutions of a polynomial system and computing the topology of an algebraic curve are fundamental problems in computational algebra, and both problems have been well-studied in the past decades. In particular, the literature on polynomial system solving is extensive, and thus, we can only provide a very brief overview, where we refer to an (incomplete) list of relevant books and papers and the references therein. We mainly distinguish between two classes of approaches for solving polynomial systems, namely, numerical and symbolic methods.

Numerical algorithms are predominantly used in engineering applications, mainly because of their high efficiency for real-world instances, their applicability for high-dimensional systems, and their relative ease of implementation. However, the major drawback of most numerical methods is that they are prone to fail in degenerate cases such as singular or clustered solutions, or if the solution set is infinite. Also, they typically do not provide any guarantees about correctness or completeness of their output. Popular representatives of the class of numerical methods are homotopy solvers such as HOM4PS by Li and Tsai [36], BERTINI by Bates et al. [49], or PHCPACK by Verschelde [52]. Their common modus operandi is to heuristically track the known solution set of a rather simple polynomial system under continuous transformations until the problem at hand is reached; e.g., consider [49, 4] for a great introduction into the field. Sophisticated strategies to ensure robustness against numerical instabilities have been developed, and recent work even introduces techniques to certify the computed solutions as so-called *approximate solutions*; e.g., see [4, 5] and the references given therein. However, although such methods can be used to compute arbitrarily good approximations of the solutions of a polynomial system, they must fail to isolate multiple roots due to the single use of approximate arithmetic. For polynomial systems with integer coefficients, a possible solution to the latter problem is to compute very good approximations of the solutions and, then, to use worst-case separation bounds for the distance between two distinct solutions. However, this approach requires extremely good approximations of the solutions for any given input, and thus, is considered to be irrelevant in practice. In addition, we are not aware of worst-case (bit) complexity bounds for such an approach that are comparable to the paper at hand.

Another kind of numerical methods is based on subdivision schemes; e.g., [1, 41, 13, 43]. Those algorithms work on a region of interest, which can a priori be chosen large enough to comprise all global solutions of the input. The region is then recursively subdivided into smaller regions until, for each of them, either the size of the region is smaller than some predefined threshold or it can be ensured that the region contains no or exactly one

unique solution. For instance, this can be achieved using predicates based on interval arithmetic or Rouché’s criterion. An important benefit is the locality of the approach, however, it requires efficient predicates for excluding empty cells to confine an exponential growth of the recursion tree, in particular in higher dimensions. In fact, for more than one dimension, we are not aware of any polynomial-time subdivision method for approximating all solutions that uses only purely numerical predicates. In addition, for isolating multiple solutions, subdivision methods share the same difficulties with other purely numerical methods. Again, for integer or rational input, worst-case separation bounds can be used to handle even degenerate cases, however, for the price of rendering the overall method impractical; e.g., see [13].

The second main approach, which is also considered in the paper at hand, rather uses *symbolic* than numerical methods for solving polynomial systems. That is, in a first step, they use elimination techniques such as the computation of resultants or Gröbner Bases to reduce the number of unknowns. For zero-dimensional systems, this means that the solutions are projected into one dimension via algebraic manipulations, followed by a computation of the projections as the roots of the corresponding elimination polynomial. Then, in a second step, the solutions are recovered from their projections. The latter (lifting) step is relatively cheap if the projection already comes along with a parametrization of the solutions as provided by a rational univariate representation (RUR); e.g., see [29, 20, 12, 45]. In contrast, the lifting step can be quite costly if the projection is not known to separate the solutions as given, at least in general, for a triangular decomposition. Computing the topology of a real algebraic plane curve $\mathcal{C} = \{(x, y) \in \mathbb{R}^2 : f(x, y) = 0\}$ can be achieved using Collin’s classical cylindrical algebraic decomposition technique [17], dating back to 1975; see also [3]. It heavily relies on the computation of the x -critical points of the curve, and thus, on the computation of the solutions of the bivariate polynomial system $f = f_y = 0$. In this context, it is remarkable that all known complete and certified algorithms, for which polynomial time bounds are known, compute (more or less) a cylindrical algebraic decomposition.

When compared to purely numerical methods, the main advantage of symbolic methods is that they can be made complete and certified, even in the presence of a degenerate situation such as multiple solutions of a polynomial system or singular points of an algebraic curve. However, this also comes with a price, namely, the high computational cost for the considered symbolic operations, which do not adapt to the actual hardness of the input. As a consequence, when solving a polynomial system, the running time mainly depends on the considered symbolic operations, and not on the geometry of the solutions. That is, it makes almost no difference whether the solutions are simple and well separated from each other or whether there exist clustered or multiple solutions. However, a considerable amount of work has been invested in the past years to reduce the number of purely symbolic operations and to replace algebraic manipulations by approximate arithmetic, without abstaining from completeness or correctness; e.g., see [50, 16, 8, 31, 14] for an (incomplete) list of recent papers that describe hybrid methods combining the computation of a cylindrical algebraic decomposition and numerical computation. As a consequence, the bounds on the theoretical complexity for the problems (P1) and (P4) have been improved over the years in an impressive manner: Arnon and McCallum [2] gave the first sub-exponential bit complexity bound $\tilde{O}(n^{30} + n^{27}\tau^3)$ for deterministically computing the topology of a planar

algebraic curve, and this bound was subsequently improved⁴ to $\tilde{O}(n^9\tau + n^8\tau^2)$ by Cheng et al. [16], Gonzalez-Vega and El Kahoui [29], Basu et al. [3, Sec. 11.6], Diochnos et al. [20], and Kerber and Sagraloff [31, 32]. A very recent manuscript⁵ even reports on a deterministic algorithm for computing the topology of an algebraic plane curve that uses $\tilde{O}(n^7 + n^6\tau)$ bit operations, and thus, lags behind our method by only one magnitude. For randomized methods, the current record bound [39] on the expected number of bit operations is even as low as $\tilde{O}(n^6 + n^5\tau)$, and thus, comparable to our bound.

For the special task of solving a bivariate polynomial systems, several projection based methods have been recently presented and analyzed. In [20], Diochnos et al. discuss three methods to solve a bivariate polynomial system. The first method, GRID, is similar to our method CBI-SOLVE and its predecessor BISOLVE in the sense that, in a first step, it determines a grid of candidate solutions via projecting the solutions onto the x - and y -axis. However, verification of the candidates is done in a completely different manner, that is, candidates are either verified as solutions or discarded by evaluating a corresponding subresultant sequence at the endpoints of isolating intervals for the projected solutions. The bit complexity of GRID is bounded by $\tilde{O}(n^{14} + n^{12}\tau^2)$, where the overall cost is dominated by that of the evaluation of the subresultant sequence. The second approach, M-RUR, is based on the computation of an RUR and achieves a bit complexity of $\tilde{O}(n^{12} + n^{10}\tau^2)$. The third approach, G-RUR, computes the greatest common divisor of the square-free parts of $f(\alpha, y)$ and $g(\alpha, y)$, where α is the projection of a solution of the system. Its bit complexity is also bounded by $\tilde{O}(n^{12} + n^{10}\tau^2)$. In [14], Cheng et al. propose the so-called local-generic position method (LGP for short). Instead of considering a coordinate transformation that is separating for all solutions, they consider, for each projection of a solution, a corresponding transformation that is locally separating for all the solutions that are located above a small neighborhood of the projection. The method shows good efficiency in practice, however, with respect to worst-case bit complexity, it suffers from the fact that such a separating form may be of a large bit-size. It is analyzed in [15], where the authors derive the bound $\tilde{O}(n^{10} + n^9\tau)$ for its bit complexity. The best current deterministic algorithms for solving bivariate systems are due to Bouzidi et al. [12, 11] and Emeliyanenko et al. [6, 24]. It has been shown that both methods isolate the solutions of a bivariate system with $\tilde{O}(n^8 + n^7\tau)$ bit operations. The method from Bouzidi et al. is based on the computation of a separating form followed by the computation of a RUR of the system. For both steps, the authors derive the bound $\tilde{O}(n^8 + n^7\tau)$ for the bit complexity. Very recent work⁶ indicates that the latter bound can be even lowered by a factor n . For randomized algorithms, the current record bound [39] on the expected number of bit operations is $\tilde{O}(n^6 + n^5\tau)$, however, the algorithm can only be used to compute the real solutions of the system $f = g = 0$ as it reduces the latter problem to the problem of computing the topology of the real algebraic curve defined by $f^2 + g^2 = 0$.

⁴The references are sorted with respect to the size of the given bound for the bit complexity of the corresponding algorithm.

⁵According to personal communication (January 2014) with Marie-Françoise Roy, who sent us a corresponding preprint of [19].

⁶According to personal communication (January 2014) with Sylvain Lazard, who sent us a preliminary version of [10] stating a deterministic worst-case bound of size $\tilde{O}(n^7 + n^6\tau)$ for the computation of a separating form for bivariate polynomial systems.

In the paper at hand, we extend and modify the algorithm BISOLVE as presented in [6] to complex system solving, and we considerably improve upon the analysis from [24]. We further show that we can solve the problems (P2) and (P3) with $\tilde{O}(n^6 + n^5\tau)$ bit operations, given that isolating regions for all complex solutions are already computed. Together with our bound for the bit complexity of CBISOLVE, this yields an improvement upon the current record bounds [12, 11] for the latter two problems by a factor n^2 and a factor n^3 , respectively. Finally, we combine CBISOLVE with the algorithm TOPNT from [8] to derive our fast method for computing the topology of an algebraic plane curve.

Notation

Throughout the paper, we use the following notations:

- Complementary to the common Landau notation $O(\cdot)$ for asymptotic behavior, $\tilde{O}(\cdot)$ means that we ignore poly-logarithmic factors, that is, $\tilde{O}(T) = O(T(\log T)^k)$, where k is any fixed integer.
- For *disks in the complex plane*, we write $D_r(m) := \{z \in \mathbb{C} : |z - m| < r\}$, where $m \in \mathbb{C}$ denotes the center and $r \in \mathbb{R}_{>0}$ the radius.
- For arbitrary complex values $w, z \in \mathbb{C}$, we define $M(z) := \max\{1, |z|\}$ and $M(z, w) := \max\{1, |z|, |w|\}$.
- We write $\log := \log_2$ for the binary logarithm, and define $\overline{\log} z := \lceil M(\log M(z)) \rceil$. That is, if $|z| \leq 2$, then $\overline{\log} z$ is 1; otherwise, $\overline{\log} z$ equals $\log |z|$ rounded up to the next integer.
- The *bitsize* $\tau(a) := \overline{\log}(|a| + 1)$ of an integer $a \in \mathbb{Z}$ is the length of its binary representation, neglecting the sign bit. Notice that $\tau(a) \geq 1$ by definition. For an arbitrary complex number $\frac{a}{b} + \frac{c}{d} \cdot i$ with integers a to d , we define its bitsize as the maximum of the values $\tau(a)$ to $\tau(d)$.
- A polynomial $F \in \mathbb{Z}[x_1, \dots, x_k]$ is of *magnitude* (n, τ) if its total degree and the bitsize of its coefficients are upper bounded by n and τ , respectively.
- For a polynomial $F \in \mathbb{C}[x]$ and an integer d with $d \geq \deg F$, we write $\text{rev}_d F := x^d \cdot F(1/x)$ for the *reverse* of $F \in \mathbb{C}[x]$ with respect to x^d , that is, the polynomial with reversed coefficient sequence of that of F considered as a d -th degree polynomial.
- For a polynomial $F \in \mathbb{C}[x]$ with pairwise distinct roots $z_1, \dots, z_k \in \mathbb{C}$, we define
 - $F^* := F / \gcd(F, F')$, the *square-free part* of F ,
 - $\text{mult}(z_i, F) := \min\{m \in \mathbb{N} : \frac{\partial^m}{\partial x^m} F(z_i) \neq 0\}$, the *multiplicity* of z_i
 - $\text{sep}(z_i, F) := \min_{j \neq i} |z_i - z_j|$, the *separation* of z_i
 - $\text{sep}(F) := \min_i \text{sep}(z_i, F)$, the *root separation* of F ,
 - $\mathcal{M}(F) := |\text{lcf}(F)| \cdot \prod_{i=1}^k M(z_i)^{\text{mult}(z_i, F)}$, the *Mahler measure* of F ,
 - $\Gamma(F) := \max_i |z_i|$, the *maximum modulus* of the roots of F ,

- $\Sigma(F) := \sum_{i=1}^k \text{mult}(z_i, f) \cdot \overline{\log}(\text{sep}(z_i, F)^{-1})$, and
- $\Sigma^*(F) := \Sigma(F^*) = \sum_{i=1}^k \overline{\log}(\text{sep}(z_i, F)^{-1})$.

If it is clear from the context, we omit F and simply write mult_i or mult_{z_i} and sep_i or sep_{z_i} in place of $\text{mult}(z_i, F)$ and $\text{sep}(z_i, F)$.

- For a square matrix $A = (a_{ij}) \in \mathbb{C}^{n \times n}$, we write $\mathcal{H}(A)$ for the row-wise Hadamard's bound on the determinant of A :

$$\mathcal{H}(A) := \left(\prod_{i=1}^n \sum_{j=1}^n |a_{ij}|^2 \right)^{1/2}.$$

2 Solving a Bivariate Polynomial System

We first consider the problem of computing isolating regions in \mathbb{C}^2 for all complex solutions of a system of equations

$$f(x, y) = \sum_{i+j \leq m} f_{i,j} x^i y^j = 0 \quad \text{and} \quad g(x, y) = \sum_{i+j \leq n} g_{i,j} x^i y^j = 0, \quad (3)$$

where $f, g \in \mathbb{Z}[x, y]$ are bivariate polynomials of total degrees $m \geq 1$ and $n \geq 1$, respectively. Furthermore, let the absolute values $|f_{i,j}|$ and $|g_{i,j}|$ of the coefficients of f and g be bounded by 2^{τ_f} and 2^{τ_g} for some positive integers $\tau_f, \tau_g \geq 1$. We assume that the solution set

$$V := \{(\alpha, \beta) \in \mathbb{C}^2 : f(\alpha, \beta) = g(\alpha, \beta) = 0\} \quad (4)$$

is *zero-dimensional* or, equivalently, that f and g share no common nontrivial factor in $\mathbb{C}[x, y] \setminus \mathbb{C}$. Then, Bézout's theorem states that the number of solutions counted with multiplicity is at most $m \cdot n$. In fact, the number of solutions equals $m \cdot n$ if we consider the corresponding homogeneous system and also count the solutions at infinity.

We assume that the input polynomials are given by their coefficients. Considering f and g as *univariate* polynomials in x and y , respectively, we write:

$$\begin{aligned} f(x, y) &= \sum_{i=0}^{m_x} f_i^{(x)}(y) x^i = \sum_{i=0}^{m_y} f_i^{(y)}(x) y^i \quad \text{and} \\ g(x, y) &= \sum_{i=0}^{n_x} g_i^{(x)}(y) x^i = \sum_{i=0}^{n_y} g_i^{(y)}(x) y^i, \end{aligned} \quad (5)$$

where $f_i^{(y)}, g_i^{(y)} \in \mathbb{Z}[x]$ and $f_i^{(x)}, g_i^{(x)} \in \mathbb{Z}[y]$, and $m_x, m_y \leq m$ and $n_x, n_y \leq n$ denote the degrees of f and g as univariate polynomials.

2.1 Some facts about resultants

Our approach for handling bivariate systems is based on elimination of variables by computing resultants. Before we present our algorithm to isolate the solutions of (3), we briefly

review some fundamental and well known results about resultants. The reader familiar with the topic may skip this part and will find pointers to the explanations in the main text. We try to keep our description as short as possible and mainly focus on our applications. For an in-depth discussion of the subject, there is an extensive range of literature for further study; e.g., we recommend [3, 28, 27, 53].

Definition 1 (Sylvester matrix). For two univariate polynomials $f = \sum_{i=0}^m f_i x^i$ and $g = \sum_{i=0}^n g_i x^i \in \mathbb{D}[x]$ of degree $m = \deg f$ and $n = \deg g$ over some integral domain \mathbb{D} , with $f \neq 0$ and $g \neq 0$, their *Sylvester matrix* $S(f, g; x) \in \mathbb{D}^{(m+n) \times (n+m)}$ with respect to x is defined as

$$S(f, g; x) := \begin{pmatrix} f_m & f_{m-1} & \cdots & f_0 & & & \\ & f_m & f_{m-1} & \cdots & f_0 & & \\ & & \ddots & & & \ddots & \\ & & & f_m & f_{m-1} & \cdots & f_0 \\ g_n & g_{n-1} & \cdots & g_0 & & & \\ & g_n & g_{n-1} & \cdots & g_0 & & \\ & & \ddots & & & \ddots & \\ & & & g_n & g_{n-1} & \cdots & g_0 \end{pmatrix}.$$

Notice that there are n rows with coefficients of f and m rows with coefficients of g .

In our application, $f, g \in \mathbb{Z}[x, y]$, and we consider both Sylvester matrices with respect to y or x as the outer variable. Hence, we write $S(f, g; y)$ for $\mathbb{D} = \mathbb{Z}[x]$ (and the entries of the Sylvester matrix are polynomials in $\mathbb{Z}[x]$), and $S(f, g; x)$ for $\mathbb{D} = \mathbb{Z}[y]$.

Definition 2 (Resultant). For f and g as in Definition 1, the *resultant* $\text{res}(f, g; x)$ of f and g with respect to x is the determinant of $S(f, g; x)$. It will be convenient to additionally define $\text{res}(0, 0; x) := 1$ and $\text{res}(f, 0; x) = \text{res}(0, g; x) := 0$ for all $f, g \neq 0$.

Theorem 3 (Properties of resultants). Let f and g be polynomials in $\mathbb{D}[x]$ as above. The resultant of f and g satisfies the following properties:

1. If $f(x) = \text{lcf } f \cdot \prod_{i=1}^m (x - x_i)$ and $g(x) = \text{lcf } g \cdot \prod_{j=1}^n (x - y_j)$, where x_i and y_j are the roots of f and g in an algebraic closure $\overline{\mathbb{D}}$ of \mathbb{D} , possibly repeated according to their multiplicities, then

$$\begin{aligned} \text{res}(f, g; x) &= (\text{lcf } f)^n \prod_i g(x_i) = (\text{lcf } g)^m \prod_j f(y_j) \\ &= (\text{lcf } f)^n (\text{lcf } g)^m \prod_{i,j} (x_i - y_j). \end{aligned}$$

2. $\text{res}(f, g; x) \equiv 0$ if and only if f and g share a common nontrivial factor in $\mathbb{D}[x] \setminus \mathbb{D}$.
3. The resultant can be represented as a $\mathbb{D}[x]$ -linear combination of f and g as

$$\text{res}(f, g; x) = u(x) \cdot f(x) + v(x) \cdot g(x),$$

where $u \in \mathbb{D}[x]$ and $v \in \mathbb{D}[x]$ are polynomials of degree less than or equal to $\deg g$ and $\deg f$, respectively. Furthermore, the polynomials u and v can be written as the determinants

of the “Sylvester-like” matrices

$$U(x) = \begin{pmatrix} f_m & \cdots & f_0 & x^{n-1} & \vdots \\ & \ddots & & \ddots & \vdots \\ & & f_m & \cdots & 1 \\ g_n & \cdots & g_0 & \cdots & 0 \\ & \ddots & & \ddots & \vdots \\ & & g_n & \cdots & 0 \end{pmatrix} \text{ and } V(x) = \begin{pmatrix} f_m & \cdots & f_0 & \cdots & 0 \\ & \ddots & & \ddots & \vdots \\ & & f_m & \cdots & 0 \\ g_n & \cdots & g_0 & x^{m-1} & \vdots \\ & \ddots & & \ddots & \vdots \\ & & g_n & \cdots & 1 \end{pmatrix},$$

which are obtained from the Sylvester matrix $S(f, g; x)$ by replacing the last column with $(x^{n-1}, \dots, 1, 0^m)$ and $(0^n, x^{m-1}, \dots, 1)$, respectively.

Let f and $g \in \mathbb{Z}[x, y]$ be bivariate polynomials without common nontrivial factor as in (3) and (5). Consider f and g as univariate polynomials over $\mathbb{D} = \mathbb{Z}[x]$, and write $R := \text{res}(f, g; y)$ for their resultant with respect to y . Then, it holds:

4. $R \in \mathbb{Z}[x]$ and $\deg R \leq m \cdot n$.
5. The roots of R are exactly the projections of the solutions of $f = g = 0$, including points at infinity, onto the complex x -plane. More precisely,

$$R(\alpha) = 0 \quad \text{if and only if} \quad \begin{cases} f(\alpha, \beta) = g(\alpha, \beta) = 0 \text{ for some } \beta \in \mathbb{C} & \text{or} \\ f_{m_y}^{(y)}(\alpha) = g_{n_y}^{(y)}(\alpha) = 0. \end{cases}$$

The multiplicity $\text{mult}(\alpha, R)$ is the sum of the intersection multiplicities⁷ of all solutions of (3) with x -coordinate α .

In order to estimate the magnitude of the resultant polynomials, we will use the following generalization of Hadamard’s bound for the size of determinants to the case of polynomial entries.

Theorem 4 (Goldstein-Graham). Let $A(x) = (a_{ij})$ be an n -square matrix whose entries $a_{ij}(x)$ are polynomials in $\mathbb{C}[x]$. Denote by $B = (b_{ij})$ the matrix with $b_{ij} = \|a_{ij}(x)\|_1$ and by $\mathcal{H}(B)$ the row-wise Hadamard bound for its determinant. The polynomial determinant $\det A(x)$ of A satisfies the inequality

$$\|\det A(x)\|_2 \leq \mathcal{H}(B).$$

Proof. See, e.g., [53, Thm. 6.31] or Lossers’ elegant solution in [38] to the claim stated by Goldstein and Graham as an exercise problem. \square

Corollary 5. If f and g as in (3) and (5) have magnitude (m, τ_f) and (n, τ_g) , respectively, then their resultant $R := \text{res}(f, g; y)$ has magnitude

$$(mn, m\tau_g + n\tau_f + O(m \log n + n \log m)).$$

⁷The multiplicity of a solution (α, β) of (3) is defined as the dimension of the localization of $\mathbb{C}[x, y]/\langle f, g \rangle$ at (α, β) , considered as a \mathbb{C} -vector space (cf. [3, Sec. 4.5, p. 148]). For a root α of $\gcd(f_{m_y}^{(y)}, g_{n_y}^{(y)})$, the intersection multiplicity at the “infinite point” (α, ∞) has also to be taken into account, however, for simplicity, we decided not to consider the more general projective setting.

Proof. The bound on the degree follows immediately from Theorem 3 (4); it remains to prove the claim on the bitsize. Choose $A = S(f, g; y)$ in Theorem 4. The 1-norms of the entries corresponding to f and g in the Sylvester matrix are bounded by $(m+1)2^{\tau_f}$ and $(n+1)2^{\tau_g}$. Since Hadamard's bound is monotone in the entries of the matrix, we can overestimate the non-zero entries of B by those values. As there are n rows with $m_y + 1 \leq m+1$ entries corresponding to f , and m rows for g , Theorem 4 yields

$$\begin{aligned} \log \|R\|_2 &\leq \log \mathcal{H}(B) = \log \left(((m+1)^{3/2} 2^{\tau_f})^n \cdot ((n+1)^{3/2} 2^{\tau_g})^m \right) \\ &= n\tau_f + m\tau_g + O(n \log m + m \log n). \end{aligned} \quad \square$$

Lemma 6. *Under the above assumptions on the magnitude of f and g , the resultant R can be computed with $\tilde{O}(\max\{m, n\}^3(m\tau_g + n\tau_f))$ bit operations.*

Proof. We only give the main ideas behind a small primes reduction and lifting approach for $R = \text{res}(f, g; y)$. For an in-depth explanation see, e.g., [23, Sec. 2.4.3 and 2.5.4]; there, Emeliyanenko discusses the symmetric case where $m = n$ and $\tau_f = \tau_g$, but the arguments carry over.⁸

Write $N = mn$ and $T = \tilde{O}(m\tau_g + n\tau_f)$ for the bounds on the degree and the bitsize of R as given in Corollary 5; here, T is an explicit bound as computed in the proof of Corollary 5. To recover the coefficients of R via Chinese remaindering, it suffices to compute the images $R_{p_i} \in \mathbb{Z}_{p_i}[x]$ of R modulo pairwise distinct primes p_1, \dots, p_r such that $\prod_i p_i > 2^T$. We choose $r = O(T)$ primes of small magnitude $\mu(p_i) = \tilde{O}(\log T)$; notice that this is possible due to the prime number theorem, and the cost is bounded by $\tilde{O}(T)$ since the cost for testing an integer of bitsize L for being prime is polynomial in T .

Using asymptotically fast methods [27, Thm. 10.24], the simultaneous reduction of any integer of bitsize bounded by $O(T)$ over all moduli p_1, \dots, p_r can be accomplished in time $\tilde{O}(T)$. Since the number of coefficients of f and g is quadratic in the total degree, reducing all coefficients requires $\tilde{O}((m^2 + n^2)T) = \tilde{O}(\max\{m, n\}^2 \cdot T)$ bit operations.

The results in [27, 44, 37] show that, for two bivariate polynomials in $\mathbb{F}[x, y]$ over a field \mathbb{F} with total degrees bounded by d , their resultant can be computed with $\tilde{O}(d^3)$ arithmetic operations in \mathbb{F} . Since the bitsizes of the p_i do not exceed $\tilde{O}(\log T)$, this translates to $\tilde{O}(\max\{m, n\}^3 \log T)$ bit operations for each individual R_{p_i} , and the total complexity of computing all R_{p_i} is $\tilde{O}(\max\{m, n\}^3 \cdot T)$.

Finally, lifting the N coefficients of R from their modular images, again using asymptotically fast methods [27, Thm. 10.25], takes $\tilde{O}(NT)$ operations. Hence, the overall time spent for the computation of R is dominated by the cost $\tilde{O}(\max\{m, n\}^3 \cdot T)$ of computing the modular resultants. \square

2.2 Root finding and multipoint evaluation

Based on an asymptotically fast method [42] for the approximate factorization of a polynomial due to Pan, recent work [39] provides an asymptotically fast method for isolating

⁸In [23, 22], Emeliyanenko reports on an extremely efficient implementation of the small primes and lifting approach on modern graphics hardware. For this, he uses a generalized Schur Algorithm for computing the determinant of a matrix with low displacement rank (as the Sylvester matrix). We remark that his implementation is also integrated in the current C++ implementation of BISOLVE [6, 8].

and approximating the roots of an arbitrary polynomial $F \in \mathbb{C}[x]$. As input the algorithm receives an oracle which can deliver arbitrarily good approximations of the coefficients, as well as the number of distinct complex roots. The following theorem summarizes the complexity results for this method; see also [46] for a dedicated real root solver that achieves a comparable bit complexity bound provided that the input polynomial is square-free.

Theorem 7. [39, Thm. 3 and Thm. 5]

1. Let $F \in \mathbb{C}[x]$ be an arbitrary polynomial of degree d with complex coefficients of absolute value less than 2^μ and leading coefficient of absolute value larger than 1. In addition, suppose that the number k of distinct roots of F is given. Then, we can compute isolating disks $D(z_i) = D_{r_i}(m_i)$ with $r_i < \text{sep}_i / (64d) = \text{sep}(z_i, F) / (64d)$ for all complex roots z_1 to z_k of F , together with their corresponding multiplicities $\text{mult}_i = \text{mult}(z_i, F)$, using

$$\tilde{O}\left(d^3 + d^2\tau + d \cdot \sum_{i=1}^k \left(\log \left(\frac{\partial^{\text{mult}_i}}{\partial x^{\text{mult}_i}} F(z_i) \right)^{-1} + \text{mult}_i \cdot \log(\text{sep}_i^{-1}) \right)\right)$$

bit operations. For this, we need an approximation of F to ρ bits after the binary point, with ρ bounded by

$$\tilde{O}\left(d \cdot \log(\Gamma(F)) + \sum_{i=1}^k \left(\log \left(\frac{\partial^{\text{mult}_i}}{\partial x^{\text{mult}_i}} F(z_i) \right)^{-1} + \text{mult}_i \cdot \log(\text{sep}_i^{-1}) \right)\right).$$

2. Let $F \in \mathbb{Z}[x]$ be a univariate polynomial of magnitude (d, μ) . Then, we can compute isolating disks $D(z_i)$ for all roots z_i as above together with their corresponding multiplicities mult_i , using a number of bit operations bounded by $\tilde{O}(d^3 + d^2\mu)$. For further refining all the isolating disks to a size of less than 2^{-L} , with L an arbitrary positive integer, we need $\tilde{O}(d^3 + d^2\mu + dL)$ bit operations.

We also need the following well-known bound for the Mahler measure of a polynomial:

Lemma 8. [3, Prop. 10.8] Let $F \in \mathbb{Z}[x]$ be a univariate polynomial of magnitude (d, μ) with pairwise distinct roots $z_1, \dots, z_m \in \mathbb{C}$. It holds that

$$\sum_{i=1}^m \text{mult}_i \cdot \log M(z_i) \leq \log \mathcal{M}(F) \leq \log \|F\|_2 = O(\mu + \log d). \quad (6)$$

The following theorem provides a bound for the sum of all values $\log \text{sep}(z_i, F)^{-1}$, where we count each value according to the multiplicity of the corresponding root. Since Theorem 9 has already been presented in [24], we decided to outsource its proof to A:

Theorem 9. Let F be as above. It holds that

$$\sum_{i=1}^m \text{mult}(z_i, F) \cdot \log \text{sep}(z_i, F)^{-1} = \tilde{O}(d^2 + d\mu).$$

In particular,

$$\Sigma^*(F) \leq \Sigma(F) = \tilde{O}(d^2 + d\mu). \quad (7)$$

Finally, we provide a complexity bound for approximately computing the values of a polynomial $F \in \mathbb{C}[x]$ at a set of $d = \deg F$ many points. The underlying algorithm essentially follows the classical fast multipoint evaluation scheme [40, 27] by Moenck and Borodin, however, the considered polynomial multiplications and divisions are carried out with approximate but certified arithmetic. We provide a self-contained proof of Theorem 10 in B.⁹

Theorem 10. *Let $F \in \mathbb{C}[x]$ be a polynomial of degree d with $\|F\|_1 \leq 2^\mu$, where $\mu \geq 1$, and let $x_1, \dots, x_m \in \mathbb{C}$ be points of absolute values bounded by 2^Γ , where $m \leq d$ and $\Gamma \geq 1$. For an arbitrary positive integer L , we can compute values $\tilde{y}_i \in \mathbb{C}$ such that $|\tilde{y}_i - F(x_i)| \leq 2^{-L}$ for all $i = 1, \dots, m$ using $\tilde{O}(d(L + \mu + d\Gamma))$ bit operations. The precision demand on the coefficients of F and the points x_i is bounded by $L + \tilde{O}(\mu + d\Gamma)$ bits after the binary point.*

Proof. See B. For alternative approaches based on a division scheme from Schönhage [47], see [35, Thm. 10] or [33, Thm. 3.9]. \square

2.3 Review of the algorithm CBISOLVE

Our bivariate system solver CBISOLVE extends the prior work BiSOLVE from [6]. Most steps in CBISOLVE are almost identical to the corresponding steps in BiSOLVE, however, it computes isolating regions for all complex solutions, whereas BiSOLVE isolates only the real solutions. In addition, CBISOLVE profits from the use of fast multipoint evaluation in the final validation phase which eventually yields a considerably improved overall complexity bound.

Because of the similarities to BiSOLVE, we give a self-contained but less detailed description of our algorithm in the paper at hand; for a deeper discussion of the key ideas of the algorithm and further details regarding an efficient implementation, the reader may consult the original paper.

We first give a brief outline of CBISOLVE, where we emphasize on the differences when compared to its predecessor. CBISOLVE consists of two main stages:

- a *projection phase*, where we project solutions onto their x - and y -coordinates using resultant computation and univariate root finding, and
- a *validation phase*, where we select the actual solutions of (3) among points in a candidate grid in \mathbb{C}^2 consisting of the preimages of the projected solutions.

We remark that BiSOLVE requires an additional intermediate *separation phase* in which the output of the projection phase is refined to obey some separation criterion. This is necessary because BiSOLVE uses a root isolation method that separates only the real roots of the corresponding resultant polynomials from each other, but the circumdisks of the isolating intervals are not allowed to contain nearby complex roots. In contrast, there is no need

⁹Kirrinnis [33] provides a bit complexity bound for the fast multipoint evaluation scheme when using approximate arithmetic. The analysis makes use of a fast method for approximate polynomial division due to Schönhage [47]. In B, we show that, as an alternative, one can also use the classical fast division scheme based on Newton iteration without diminishing the complexity results. Since the bounds provided by Kirrinnis are stated in a slightly different form (e.g., the polynomials have to rescaled) and since the paper seems to be not well known in the context of multipoint evaluation, we decided to give general bounds, which can directly be used in our analysis of the algorithm CBISOLVE.

for an additional separation step in CBISOLVE as we use a complex root isolator in the projection phase and, thus, the corresponding requirements are automatically satisfied. For achieving good complexity bounds, we propose to use the complex root isolator from [39] for the first step. However, in practice, it can be replaced by any certified root solver, and we suggest to use numerical methods like [9, 34] instead. We now give details:

Projection phase The initial stage is to compute isolating disks in \mathbb{C} for the x - and y -values of the solutions of (3). We define

$$\begin{aligned} V^{(x)} &:= \{\alpha \in \mathbb{C} : \exists y \in \mathbb{C} \cup \{\infty\} \text{ with } f(\alpha, y) = g(\alpha, y) = 0\} \quad \text{and} \\ V^{(y)} &:= \{\beta \in \mathbb{C} : \exists x \in \mathbb{C} \cup \{\infty\} \text{ with } f(x, \beta) = g(x, \beta) = 0\}, \end{aligned}$$

where $f(\alpha, \infty) = g(\alpha, \infty) = 0$ means that the leading coefficients $f_{m_y}^{(y)}$ and $g_{n_y}^{(y)} \in \mathbb{C}[x]$ of f and g , considered as polynomials in $(\mathbb{C}[x])[y]$, share the common factor $(x - \alpha)$.

The set V of solutions as defined in (4) is a subset of the Cartesian product

$$\mathcal{C} := V^{(x)} \times V^{(y)} \subset \mathbb{C}^2, \quad (8)$$

which we denote the *set of candidate solutions* for (3). Theorem 3 (5) states that $V^{(x)}$ is exactly the set of roots of the resultant $R^{(y)} := \text{res}(f, g; y)$ and, analogously, $V^{(y)} = \{\beta \in \mathbb{C} : R^{(x)}(\beta) := \text{res}(f, g; x)(\beta) = 0\}$. We represent the zeros of $R \in \{R^{(y)}, R^{(x)}\}$ by a sets of *isolating disks* $D^{(x)}(\alpha) := D_{r^{(x)}(\alpha)}(m^{(x)}(\alpha))$ and $D^{(y)}(\beta) := D_{r^{(y)}(\beta)}(m^{(y)}(\beta))$. For the sake of readability, we omit the directions of the projections from now and write $D(\alpha)$ for $D^{(x)}(\alpha)$ or $r(\beta)$ for $r^{(y)}(\beta)$ and so on.

We say that the set of disks is *isolating* if the (closed) disks $D(\alpha)$ and $D(\alpha')$ are disjoint for any two distinct roots α and α' . Note that, in general, $\alpha \neq m(\alpha)$. However, there is a one-to-one correspondence between the set of isolating disks $D(\alpha)$ (and, thus, their centers $m(\alpha)$) and the roots α . If the context is unambiguous, we exploit this correspondence and write α instead of $D(\alpha)$, hiding the fact that we only compute the isolating disks.

We additionally impose the following “*well-isolation*” requirements on our isolating disks. They ensure that some neighborhood of the boundaries of the disks $D(\alpha)$ contains no root, and that the disks are not needlessly small. The reason for these restrictions will become apparent in the next paragraph.

- (Isol 1) Each α is contained in the concentric disk $D_{r(\alpha)/2}(m(\alpha))$ of halved radius;
- (Isol 2) two distinct isolating disks $D(\alpha)$ and $D(\alpha')$ are separated by a margin of at least $2 \max\{r(\alpha), r(\alpha')\}$; and
- (Isol 3) the size of each disk is related to the separation and the absolute value of the corresponding root. More precisely,

$$\frac{1}{32} \min\{\text{sep}_\alpha, M(\alpha)\} \leq r(\alpha) \leq \frac{1}{4} \min\{\text{sep}_\alpha, M(\alpha)\},$$

where $r(\alpha)$ is a power of two, and the precision of the center $m(\alpha)$ matches the accuracy of the radius (up to an additive constant).

Given such well-isolating disks, each solution of the initial system (3) is contained in a *candidate polydisk* $\Delta(\alpha, \beta) := D(\alpha) \times D(\beta) \subset \mathbb{C}^2$, representing a candidate point $\xi := (\alpha, \beta) \in \mathcal{C}$ in (8), and each of these polydisks contains at most one solution (namely ξ , if and only if ξ is a solution of (3)).

Computing disks which obey the above restrictions from the output of any univariate root complex solver is straightforward: For a root α of R , let $m_0(\alpha)$ and $r_0(\alpha)$ denote the center and the radius of the isolating disks returned by the solver, respectively. Assume that $r_0(\alpha) \leq \text{sep}_\alpha / 32$ for all α , a property that is typically ensured by design of the solver; otherwise, refine the disks further. We remark that the solver presented in [39] returns such disks together with the corresponding multiplicities of the isolated roots by default. For an arbitrary but fixed root α of R , let $\alpha^* := \arg \min_{\alpha'} |m_0(\alpha) - m_0(\alpha')|$ denote a root of R whose isolating disk's midpoint is closest to $m_0(\alpha)$. Then,

$$|m_0(\alpha) - m_0(\alpha^*)| - r_0(\alpha) - r_0(\alpha^*) \leq \text{sep}_\alpha \leq |m_0(\alpha) - m_0(\alpha^*)| + r_0(\alpha) + r_0(\alpha^*)$$

and, since both sep_α and sep_{α^*} are upper bounded by $|\alpha - \alpha^*|$,

$$\text{sep}_\alpha \in \left[\frac{29}{32}, \frac{35}{32} \right] \cdot |m_0(\alpha) - m_0(\alpha^*)|.$$

Hence,

$$r(\alpha) := \frac{1}{8} \cdot 2^{\lceil \log \min\{M(\alpha), |m_0(\alpha) - m_0(\alpha^*)|\} \rceil}$$

fulfills the inequality in (Isol 3).

We can now approximate the root α by a complex number $m(\alpha)$ with dyadic real and imaginary parts such that $|m(\alpha) - \alpha| \leq r(\alpha)$. W.l.o.g., we can assume that $m(\alpha)$ is chosen in a manner such that it has no more than $1 + \lceil \log(r(\alpha)^{-1}) \rceil$ bits after the binary point. Then, we define $D(\alpha) := D_{r(\alpha)}(m(\alpha))$, a disk that satisfies (Isol 1)–(Isol 3).

Validation phase It remains to select the candidates that actually contribute to the solution set V and to discard the remaining ones. As an *exclusion predicate*, we simply use interval arithmetic with gradually increasing precision: if $\xi = (\alpha, \beta)$ is not a solution, sufficiently accurate refinement of the candidate polydisk and interval evaluation will eventually reveal that either $0 \notin f(\Delta(\xi))$ or $0 \notin g(\Delta(\xi))$. However, proper solutions cannot be certified by interval arithmetic alone. We overcome this problem by a per-candidate sandwich bound argument on the resultant values. In what follows, we write $n^* := \max\{m, n\}$.

Recall that, due to the identity from Theorem 3 (3), we can represent the resultants as

$$\begin{aligned} R^{(y)}(x) &= u^{(y)}(x, y) \cdot f(x, y) + v^{(y)}(x, y) \cdot g(x, y) \quad \text{and} \\ R^{(x)}(y) &= u^{(x)}(x, y) \cdot f(x, y) + v^{(x)}(x, y) \cdot g(x, y), \end{aligned}$$

where the cofactor polynomials $u^{(*)}$ and $v^{(*)}$ are the determinants of Sylvester-like matrices $U^{(*)}$ and $V^{(*)}$. Suppose that $\text{LB}(\alpha)$ and $\text{LB}(\beta)$ are positive lower bounds for the absolute values of $R^{(y)}$ and $R^{(x)}$, respectively, when restricted to the boundaries of $D(\alpha)$ and $D(\beta)$. In addition, suppose that $\text{UB}_w(\xi)$ is an upper bound for the magnitude of w when restricted

to $\Delta(\xi)$, with $w \in \{u^{(y)}, v^{(y)}, u^{(x)}, v^{(x)}\}$. Then, a homotopy argument [8, Thm. 4] shows that at least one of the following inequalities holds for all $(x, y) \in \Delta(\xi)$ unless ξ solves the initial system (3):

$$\begin{aligned} \text{UB}_{u^{(y)}}(\xi) \cdot |f(x, y)| + \text{UB}_{v^{(y)}}(\xi) \cdot |g(x, y)| &\geq \text{LB}(\alpha) \quad \text{or} \\ \text{UB}_{u^{(x)}}(\xi) \cdot |f(x, y)| + \text{UB}_{v^{(x)}}(\xi) \cdot |g(x, y)| &\geq \text{LB}(\beta). \end{aligned} \quad (9)$$

Obviously, the converse holds as well: If ξ is a solution, then f and g converge to zero in a neighborhood of ξ . Hence, each point $\xi_0 \in \Delta(\xi)$ that is sufficiently close to ξ violates (9), and thus, serves as a certificate for ξ being a solution. In other words, (9) serves as an *inclusion predicate* for the candidates. It remains to show how to compute such lower and upper bounds $\text{LB}(\alpha)$ and $\text{UB}(\alpha)$.

The aforementioned “well-isolating” properties (Isol 1) and (Isol 2) of an isolating disk guarantee that $\text{LB}(\alpha) := 2^{\text{lb}(\alpha)}$ and $\text{LB}(\beta) := 2^{\text{lb}(\beta)}$, with

$$\begin{aligned} \text{lb}(\alpha) &:= \left\lfloor \log |R^{(y)}(m(\alpha) - r(\alpha))| \right\rfloor - \text{mult}(\alpha, R^{(y)}) - \deg R^{(y)} \quad \text{and} \\ \text{lb}(\beta) &:= \left\lfloor \log |R^{(x)}(m(\beta) - r(\beta))| \right\rfloor - \text{mult}(\beta, R^{(x)}) - \deg R^{(x)}, \end{aligned} \quad (10)$$

define lower bounds on the magnitude of the resultants on the boundary of $D(\alpha)$ and $D(\beta)$. For a simple proof of the latter claim, see Lemma 15 in A.

For the upper bounds UB_w for the cofactor polynomials, we define

$$\begin{aligned} \text{ub}(\alpha) &:= \overline{\log} \left((3M(\alpha))^{n^*} \cdot ((m+1)^{2\tau_f} (3M(\alpha))^m)^n \cdot ((n+1)^{2\tau_g} (3M(\alpha))^n)^m \right) \\ &= \tilde{O}(mn \overline{\log} \alpha + n\tau_f + m\tau_g), \\ \text{ub}(\beta) &:= \overline{\log} \left((3M(\beta))^{n^*} \cdot ((m+1)^{2\tau_f} (3M(\beta))^m)^n \cdot ((n+1)^{2\tau_g} (3M(\beta))^n)^m \right) \\ &= \tilde{O}(mn \overline{\log} \beta + n\tau_f + m\tau_g), \\ \text{UB}(\alpha) &:= 2^{\text{ub}(\alpha)}, \quad \text{UB}(\beta) := 2^{\text{ub}(\beta)}, \quad \text{and} \\ \text{UB}(\xi) &:= \text{UB}(\alpha) \cdot \text{UB}(\beta) = 2^{\text{ub}(\alpha)} \cdot 2^{\text{ub}(\beta)}. \end{aligned} \quad (11)$$

Using Laplace expansion along the last column and the row-wise Hadamard’s determinant bound from Theorem 4 on the remaining minors of $U^{(*)}$ and $V^{(*)}$, it is straightforward to verify that the value $\text{UB}(\xi)$ constitutes an upper bound for the absolute values of the cofactors $u^{(*)}$ and $v^{(*)}$, restricted to $\Delta(\xi)$.

Notice that neither $\text{LB}(\alpha)$, $\text{LB}(\beta)$ nor $\text{UB}(\xi)$ solely depend on the exact values of α and β or the cofactor $w \in \{u^{(y)}, v^{(y)}, u^{(x)}, v^{(x)}\}$, but rather on the inclusion disks as computed in the projection phase.

It should be remarked that, in practice, it is not advisable to compute the values $\text{lb}(\alpha)$ and $\text{ub}(\alpha)$ exactly. Instead, it suffices to compute constant-factor approximations to the resultant values $R^{(*)}(m(\alpha) - r(\alpha))$ and the magnitudes $M(\alpha)$ and $M(\beta)$ and, hence, to conservatively under- and overestimate $\text{lb}(\alpha)$ and $\text{ub}(\alpha)$ by a small additive constant. Our analysis in the following section will show that this will not harm the complexity results; however, we keep the above definitions for the sake of readability.

In order to certify or discard the candidates, we could use a straightforward method based on interval arithmetic to certify or discard a candidate $\xi = (\alpha, \beta)$. That is, we iteratively refine the corresponding inclusion polydisk $\Delta(\xi)$ in stages with exponentially increasing precision (say, $\rho = 1, 2, 4, 8, 16, \dots$ bits) and use interval arithmetic to evaluate f and g at $\Delta(\xi)$. In each iteration, we check whether interval arithmetic yields that $0 \notin f(\Delta(\xi))$ or $0 \notin g(\Delta(\xi))$ or, otherwise, whether there exists a point $\xi_0 \in \Delta(\xi)$ that violates (9). Termination and correctness of the procedure follows from the above considerations.

The reader may notice that we could use fast approximate multipoint evaluation to group $O(n)$ many of such evaluations together. This is due to a natural correspondence between approximate evaluation and interval arithmetic: if an input precision of ρ bits after the binary point allows us to approximate the value of the polynomial $F \in \{f, g\}$ at the point ξ by some \tilde{F} with $|F(\xi) - \tilde{F}| < 2^{-L}$, then $|F(x) - \tilde{F}| < 2^{-L}$ for all x with $\|x - \xi\|_\infty < 2^{-\rho}$. Thus, it is not surprising that fast approximate multipoint evaluation techniques could save a factor n with respect to the computational complexity.

However, we can even do better by considering a more refined evaluation scheme: Informally speaking, the worst-case complexity of a candidate is determined by the minimum of the ratios $\text{LB}(\alpha)/\text{UB}(\alpha)$ and $\text{LB}(\beta)/\text{UB}(\beta)$. If, say, the former ratio is very small for a specific value α^* , we have to expect a high precision demand until we can decide a candidate (α^*, β) irregardless of the fact that β might be “nice”. As a consequence, when carrying out all evaluations in a naive manner, we would have to approximate the coefficients of each polynomial $f(\alpha, y)$ (and $g(\alpha, y)$) to the worst-case precision for each given fiber. Then, $O(\max\{m, n\})$ multipoint evaluations to the same precision are required to decide the $O(mn)$ candidates over the fiber $x = \alpha$ in chunks of m or n candidates, respectively, for $f(\alpha, y)$ and $g(\alpha, y)$.

In contrast, we propose to first partition the candidates and then to apply Theorem 10 when using multipoint evaluation. We first define values

$$\begin{aligned} \rho(\alpha) &:= \lceil \log(\text{UB}(\alpha)/\text{LB}(\alpha)) \rceil \quad \text{and} \\ \rho(\beta) &:= \lceil \log(\text{UB}(\beta)/\text{LB}(\beta)) \rceil, \end{aligned} \tag{12}$$

which induce a partition of the set of candidates $(\alpha, \beta) \in \mathcal{C}$ into sets

$$\begin{aligned} \mathcal{C}^{(x)} &:= \{(\alpha, \beta) \in \mathcal{C} : \rho(\alpha) \geq \rho(\beta)\} \quad \text{and} \\ \mathcal{C}^{(y)} &:= \{(\alpha, \beta) \in \mathcal{C} : \rho(\alpha) < \rho(\beta)\} = \mathcal{C} \setminus \mathcal{C}^{(x)}, \end{aligned}$$

Then, for each candidate in the first set, the precision needed to process the candidate is directly related to the “nicer” ratio $\text{LB}(\alpha)/\text{UB}(\alpha)$ corresponding to the x -coordinate of the candidate. Vice versa, for each candidate in the second set, the required precision is bounded in terms of the ratio $\text{LB}(\beta)/\text{UB}(\beta)$ corresponding to the y -coordinate.

For an arbitrary but fixed α , the candidates $(\alpha, \beta) \in \mathcal{C}^{(x)}$ are now processed in rounds as follows: Initially, we define $\rho := 1$ as the initial precision and $\mathcal{A} := \{(\alpha, \beta) \in \mathcal{C}^{(x)}\}$ as the list of active candidates. In each round, we compute approximations of the polynomial coefficients $f_i^{(y)}(\alpha)$ and $g_i^{(y)}(\alpha)$ of $f(\alpha, y)$ and $g(\alpha, y)$, respectively. Then, we use fast approximate multipoint evaluation to approximate the values $f(\alpha, \beta)$ and $g(\alpha, \beta)$, where we consider blocks of at most m and at most n distinct values β with $(\alpha, \beta) \in \mathcal{A}$. The precision for the considered evaluations is chosen such that an absolute output precision of ρ bits after the

binary point can be guaranteed. We remark that there is no need to determine this precision a priori: it suffices to iteratively increase the working precision until the output precision of ρ is achieved. If we double the precision in each step, the total cost is dominated (up to logarithmic factors) by the cost of the last run.

Candidates for which we can show that $|f(\alpha, \beta)| \geq 2^{-\rho}$ or $|g(\alpha, \beta)| \geq 2^{-\rho}$, implying $f(\alpha, \beta) \neq 0$ or $g(\alpha, \beta) \neq 0$, are discarded; candidates that violate the inequalities in (9) are stored as solutions. Each candidate that is either discarded or stored as a solution is removed from \mathcal{A} . We then continue with the next round, where we double the precision ρ . We stop if \mathcal{A} becomes empty. Following this approach, we can process all candidates $(\alpha, \beta) \in \mathcal{C}^{(x)}$. The candidate set $\mathcal{C}^{(y)}$ is handled in almost the same manner, however, we aggregate evaluations along horizontal fibers $x = \beta$. That is, we first evaluate the coefficients $f_i^{(x)}(y)$ and $g_i^{(x)}(y)$ of f and g (considered as polynomials in x) at the values $y = \beta$, and then use multipoint evaluation along the horizontal fibers.

2.4 Complexity analysis of CBISOLVE

We write N and T , with $N = m \cdot n$ and $T = m\tau_g + n\tau_f + O(m \log n + n \log m)$, for the explicit bounds on the degree and the bitsize of the resultants $R^{(y)}$ and $R^{(x)}$, as determined in Corollary 5. In addition, we define $n^* := \max\{m, n\}$ and $\tau^* := \max\{\tau_f, \tau_g\}$.

Projection phase According to Lemma 6, computing the resultants $R \in \{R^{(y)}, R^{(x)}\}$ requires $\tilde{O}(n^{*3}T)$ bit operations. For isolating the roots of R , we use the algorithm from [39, Thm. 5], which isolates the roots of R in time $\tilde{O}(N^3 + N^2T)$; cf. Theorem 7. Post-processing its output to satisfy the well-isolating properties (Isol 1)–(Isol 3) is possible within the same complexity bound; in fact, we can compute isolating disks of size 2^{-L} in time $\tilde{O}(N^3 + N^2T + NL)$. From the well isolating properties, we further conclude that the midpoints and the radii of the disks $D(\alpha)$ and $D(\beta)$ are dyadic numbers representable by $\tilde{O}(N^2 + NT)$ bits; cf. Theorem 9.

Validation phase We now analyze the worst-case precision for which success of either the inclusion test or the exclusion test for a candidate $\xi := (\alpha, \beta)$ is guaranteed. Define

$$\delta(\xi) := \min\{\text{LB}(\alpha) / \text{UB}(\alpha), \text{LB}(\beta) / \text{UB}(\beta)\}.$$

If ξ is a solution, (9) is violated if and only if there exists a point $(x_0, y_0) \in \Delta(\xi)$ with

$$|f(x_0, y_0)| + |g(x_0, y_0)| < \delta(\xi). \quad (13)$$

By contraposition, $|f(x, y)| + |g(x, y)| \geq \delta(\xi)$ for all $(x, y) \in \Delta(\xi)$ if ξ is not a solution. Hence, in order to certify or discard the candidate ξ , it suffices to approximate $f(\xi)$ as well as $g(\xi)$ to an error of less than $\delta(\xi)/2$. Namely, we can then either verify that at least one of the polynomials f and g does not vanish at ξ , or verify that (13) holds. From the definition (12) of the values $\rho(\alpha)$ and $\rho(\beta)$ and the definition of the sets $\mathcal{C}^{(x)}$ and $\mathcal{C}^{(y)}$, we conclude that each candidate $\xi = (\alpha, \beta) \in \mathcal{C}^{(x)}$ (or $\xi \in \mathcal{C}^{(y)}$) can either be discarded or verified as a solution if we approximate $f(\xi)$ as well as $g(\xi)$ to an error of less than $2^{-\rho(\alpha)-1}$ (or $2^{-\rho(\beta)-1}$), since $\rho(\alpha) \geq \rho(\beta)$ (or $\rho(\alpha) < \rho(\beta)$).

According to Theorem 10, this implies that our inclusion/exclusion predicate succeeds for a candidate $\xi = (\alpha, \beta) \in \mathcal{C}^{(x)}$ if we approximate the corresponding coefficient polynomials and run the multipoint evaluation with an absolute precision $\rho(\xi)$ of size

$$\rho(\xi) = \rho(\alpha) + 1 + \tilde{O}(\tau^* + n^* \log \alpha + n^* \log \beta) = \tilde{O}(\rho(\alpha)), \quad (14)$$

where we used that the coefficients of $f(\alpha, y)$ and $g(\alpha, y)$ have a bitsize bounded by $O(\tau^* + n^* \cdot \log \alpha)$. To see the second bound, examine the definitions of $\text{ub}(\alpha)$ and $\text{ub}(\beta)$ and verify that $n^* \log \beta \leq \rho(\beta) \leq \rho(\alpha)$ and $n^* \log \alpha \leq \rho(\alpha)$ for all α and β .

In order to carry out the above evaluations, it suffices to approximate each candidate $\xi \in \mathcal{C}^{(x)}$ to at most $\tilde{O}(\max_{\alpha} \rho(\alpha))$ bits after the binary point. This can be achieved simultaneously for all candidates by approximating all roots of the resultant polynomials $R^{(x)}$ and $R^{(y)}$ to a corresponding number of bits after the binary point, a computation that uses $\tilde{O}(N^3 + N^2T + N \max_{\alpha} \rho(\alpha))$ bit operations according to Theorem 7, Part 2.

For an arbitrary but fixed α , the cost for computing sufficiently good approximations of the polynomials $f(\alpha, y)$ and $g(\alpha, y)$ is bounded by

$$\tilde{O}(n^*(n^* \rho(\alpha) + \tau^* + n^* \log M(\alpha))) = \tilde{O}(n^{*2} \rho(\alpha)) \quad (15)$$

bit operations, and the same bound also holds for the cost of evaluating $f(\alpha, y)$ and $g(\alpha, y)$ at the $O(mn)$ many points $y = \beta$ for which $\rho(\alpha) \geq \rho(\beta)$; cf. Theorem 10. Notice that, for approximating the polynomials $f(\alpha, y)$ and $g(\alpha, y)$, the factor n^* in (15) is due to the fact that we have to evaluate up to n^* many coefficient polynomials. In contrast, for approximating the values $f(\alpha, \beta)$ and $g(\alpha, \beta)$, the factor n^* is related to the number of blocks (each consisting of at most n^* many points $y = \beta$) in the multipoint evaluation. Now, summing up the bit complexity bound in (15) over all α and adding the cost for computing sufficiently good approximations of the candidates yields the following bound for the cost of processing all candidates $(\alpha, \beta) \in \mathcal{C}^{(x)}$:

$$\tilde{O}\left(N^3 + N^2T + N \cdot \max_{\alpha} \rho(\alpha) + n^{*2} \cdot \sum_{\alpha} \rho(\alpha)\right). \quad (16)$$

It remains to exploit amortization effects for the sum

$$\sum_{\alpha} \rho(\alpha) \leq \sum_{\alpha} \log \text{UB}(\alpha) + \sum_{\alpha} \log (\text{LB}(\alpha)^{-1}).$$

From the definition of $\text{UB}(\alpha)$, we conclude that each term in the first sum on the right hand side is upper bounded by $\tilde{O}(N \log \alpha + T)$, and thus, the overall sum is upper bounded by $\tilde{O}(NT)$ according to Lemma 8 and the fact that there exist at most N many values α . For the sum related to the lower bounds, the considerations in [24, Sec. 4.3.1] carry over¹⁰ and yield

$$\begin{aligned} \sum_{\alpha} \log (\text{LB}(\alpha)^{-1}) &= \tilde{O}\left(\Sigma(R) + NT + \mathcal{M}(R) + N \sum_{\alpha} \log \alpha\right) \\ &= \tilde{O}(N^2 + NT); \end{aligned}$$

for the sake of completeness, we provide a proof in Lemma 16 in A. We conclude that $\sum_{\alpha} \rho(\alpha) = \tilde{O}(N^2 + NT)$.

¹⁰In [24], the authors discuss the symmetric case where $n = m$ and $\tau = \tau_f = \tau_g$. However, their analysis only involves the magnitude of the resultants, hence the result stated here follows directly from substituting N and T for n^2 and $n\tau$, respectively.

Substituting the latter bound in (16) shows that processing all candidates $(\alpha, \beta) \in \mathcal{C}^{(x)}$ needs $\tilde{O}(n^{*2}(N^2 + NT))$ bit operations. Finally, an entirely analogous argument as for $\mathcal{C}^{(x)}$ further shows that the latter bound also holds for the cost of processing the candidates in $\mathcal{C}^{(y)}$. We summarize:

Theorem 11. *The algorithm CBISOLVE computes isolating polydisks for all complex solutions of the system $f(x, y) = g(x, y) = 0$, where $f, g \in \mathbb{Z}[x, y]$ are bivariate polynomials of magnitude (m, τ_f) and (n, τ_g) as defined as in (3), in a total number of*

$$\tilde{O}(\max^2\{m, n\} \cdot (m^2n^2 + mn(m\tau_g + n\tau_f)))$$

bit operations. For refining all polydisks to a size of less than 2^{-L} , with L an arbitrary given positive integer, it needs

$$\tilde{O}(\max^2\{m, n\} \cdot (m^2n^2 + mn(m\tau_g + n\tau_f)) + mn \cdot L)$$

bit operations.

2.5 Computation of a separating form

In this section, we consider the problem of computing a *separating form* for the polynomial system (3), that is, a polynomial $l_s(x, y) = x + s \cdot y$ such that $\alpha + s \cdot \beta \neq \alpha' + s \cdot \beta'$ for each pair (α, β) and (α', β') of distinct solutions of (3). Typically, in most approaches for computing the solutions of a polynomial system, such an l_s is determined first, followed by a shearing $(x, y) \mapsto (x + s \cdot y, y)$ to put the system into generic position (i.e. no two solutions share the same x -coordinate). In contrast, we assume that all solutions $\xi_i := (x_i, y_i) \in \mathbb{C}$, $i = 1, \dots, r$ with some $r \leq m \cdot n$, of the input system are already computed using our algorithm CBISOLVE and derive a separating form from sufficiently good approximations of the solutions. This seems to be artificial at first glance since, usually, the main reason for computing a separating form is to solve the system. However, in Section 3 we will show that we can derive the topology of the algebraic curve defined as the zero set of $f(x, y)$ from the solutions of the systems $f = f_y = 0$ and $f_x = f_y = 0$ and a corresponding separating form for the latter system.

Since there exist at most $m \cdot n$ distinct solutions, there can be at most $\binom{m \cdot n}{2}$ values for s yielding a linear form that is not separating. Namely, all such "bad" values for s must be among the set of all values

$$s_{ij} := \frac{x_i - x_j}{y_j - y_i} \in \mathbb{C}, \quad \text{where } i, j \in \{1, \dots, n\} \text{ and } y_i \neq y_j. \quad (17)$$

Hence, in order to compute a separating form, it suffices to approximate each value s_{ij} by some $\tilde{s}_{ij} \in \mathbb{C}$ with $|s_{ij} - \tilde{s}_{ij}| < 1/2$ and to choose an integer $s \in \{0, \dots, \binom{m \cdot n}{2}\}$ which is not contained in any of the disks $D_{1/2}(\tilde{s}_{ij})$. Namely, following this approach, we can exclude at most one integer value for s from each pair of solutions, and thus, there is at least one s left from the remaining integers that yields a separating form.

We estimate the cost for computing a separating form in the above described way: For approximating s_{ij} to an absolute error of $1/2$, it suffices to approximate $x_i - x_j$ as well as $y_i - y_j$ to a number of $L_{ij} = O(\log \max(|x_i - x_j|, |y_i - y_j|) + \log((y_i - y_j)^{-1}))$ bits after the

binary point and to divide the two values using approximate arithmetic with a precision of L_{ij} . If we assume that corresponding approximations of the solutions (x_i, y_i) and (x_j, y_j) are already given, then the cost for carrying out the two subtractions and the division is bounded by $\tilde{O}(L_{ij})$. Again, note there is no need to calculate the working precision a priori: We can carry out the evaluations with an absolute precision $L = 1, 2, 4, 8, \dots$ until we can approximate \tilde{s}_{ij} to an absolute error of $1/2$, and we are guaranteed to succeed for some L of size $O(\log \max(|x_i - x_j|, |y_i - y_j|) + \log((y_i - y_j)^{-1}))$. The total cost is dominated by the last call.

In the previous section, we have shown that the pairwise distance between any two distinct values x_i and x_j as well as between any two distinct values y_i and y_j is lower bounded by $2^{-\tilde{O}(N^2 + NT)}$, where we again write $N = mn$ and T , with $T = \tilde{O}(m\tau_g + n\tau_f)$, for the bounds on the degree and the bitsize of the resultants of f and g . Furthermore, the absolute value of all x_i and all y_i is upper bounded by $2^{O(T)}$. Hence, it follows that $L_{ij} = \tilde{O}(NT)$ for all pairs (i, j) , and thus, the solutions (x_i, y_i) have to be approximated to $\tilde{O}(NT)$ bits after the binary points. The cost for computing such approximations of all solutions is bounded by $\tilde{O}(n^{*2}(N^2 + NT))$ according to Theorem 11, where $n^* := \max\{m, n\}$. It remains to bound the cost for the evaluations needed to compute the values \tilde{s}_{ij} . For this, we fix some index i and sum up the precisions L_{ij} over all j with $y_j \neq y_i$:

$$\begin{aligned} \sum_{j: y_j \neq y_i} L_{ij} &= O\left(\sum_{j: y_j \neq y_i} \log \max(|x_i - x_j|, |y_i - y_j|) + \sum_{j: y_j \neq y_i} \log((y_i - y_j)^{-1})\right) \\ &= \tilde{O}\left(NT + \sum_{\beta \neq y_i: R^{(x)}(\beta)=0} \text{mult}(\beta, R^{(x)}) \cdot \log((y_i - \beta)^{-1})\right) \\ &= \tilde{O}\left(NT + \sum_{\beta \neq y_i: R^{(x)}(\beta)=0} \text{mult}(\beta, R^{(x)}) \cdot \log(\text{sep}(\beta, R^{(x)})^{-1})\right) \\ &= \tilde{O}(N^2 + NT), \end{aligned}$$

where $R^{(x)} := \text{res}(f, g; x)$ and β runs over all distinct roots of $R^{(x)}$ that are different from y_i . For the second inequality, we used that each y_j is a root of $R^{(x)}$ and that there are at most $\text{mult}(\beta, R^{(x)})$ many pairs (x_j, y_j) with $\beta = y_j$; namely, each such solution contributes with at least one to the multiplicity of β . For the last inequality, we used that, for an arbitrary integer polynomial F of magnitude (d, μ) , we have $\sum_{z: F(z)=0} \text{mult}(z, F) \cdot \log(\text{sep}(z, F)^{-1}) = \tilde{O}(d^2 + d \cdot \mu)$; see Theorem 9. It follows that the sum over all L_{ij} with $y_i \neq y_j$ is bounded by $\tilde{O}(N^3 + N^2T)$ as there are at most N solutions (x_i, y_i) . We summarize:

Theorem 12. *A separating form $l_s(x, y) = x + s \cdot y$ with $s \in \{0, \dots, \binom{m+n}{2}\}$ for the polynomial system (3) can be computed with a number of bit operations bounded by*

$$\tilde{O}(\max\{m, n\} \cdot (m^2n^2 + mn(m\tau_g + n\tau_f))).$$

2.6 Sign evaluation of a polynomial at the real solutions

We next study the problem of evaluating the sign of a polynomial $h \in \mathbb{Z}[x, y]$ at the *real valued* solutions of the system (3). In order to simplify the presentation, we assume throughout the following considerations that f, g , and h are integer polynomials of magnitude (n, τ) .

We first consider the case, where h shares only a non-trivial factor with at least one of the polynomials f and g . Then, w.l.o.g., we can assume that $\gcd(g, h) \in \mathbb{Z} \setminus \{0\}$; otherwise, switch f and g . According to Theorem 11, we can use CBISOLVE to compute isolating polydisks in \mathbb{C}^2 for all solutions of $f = g = 0$ as well as for $g = h = 0$ with $\tilde{O}(n^6 + n^5\tau)$ bit operations. Recall that CBISOLVE also computes the resultant polynomials $R^{(y)} = \text{res}(f, g; y) \in \mathbb{Z}[x]$ and $R^{(x)} = \text{res}(f, g; x) \in \mathbb{Z}[y]$ with $R^{(y)}(\alpha) = R^{(x)}(\beta) = 0$ and corresponding isolating (and refineable) disks $D(\alpha)$ and $D(\beta)$ for the coordinates of each solution (α, β) of $f = g = 0$. For the system $g = h = 0$, it computes corresponding integer polynomials $\bar{R}^{(y)} := \text{res}(g, h; y)$ and $\bar{R}^{(x)} := \text{res}(g, h; x)$ with $\bar{R}^{(y)}(\bar{\alpha}) = \bar{R}^{(x)}(\bar{\beta}) = 0$ and corresponding isolating polydisks $D(\bar{\alpha})$ and $D(\bar{\beta})$ for the coordinates of the solutions $(\bar{\alpha}, \bar{\beta})$ of $g = h = 0$. Now, in order to determine the common solutions of $f = g = 0$ and $g = h = 0$, we can simply compare the roots of the polynomials $R^{(x)}$ and $R^{(y)}$ with those of the polynomials $\bar{R}^{(x)}$ and $\bar{R}^{(y)}$, respectively. In [39, Lem. 15], it has been shown that the cost for comparing the roots is bounded by $\tilde{O}(n^6 + n^5\tau)$ bit operations.¹¹ It remains to evaluate the sign of h at those real valued solutions of $f = g = 0$ that are *not* solutions of $g = h = 0$. This can be achieved in a straightforward manner using approximate evaluation. More precisely, we approximate $h(\alpha, \beta)$ to a precision of $L = 1, 2, 4, 8, \dots$ bits after the binary point. We stop increasing L as soon as we can decide the sign of $h(\alpha, \beta)$. This is the case if the approximation of $h(\alpha, \beta)$ has absolute value larger than 2^{-L} . The cost for determining the sign of $h(\alpha, \beta)$ is dominated by the call for the target precision and, thus, bounded by $\tilde{O}(n^2(\tau + \log(h(\alpha, \beta)^{-1}) + n \cdot \log M(\alpha, \beta)))$ bit operations.

In order to bound the overall cost for the sign evaluations, we derive upper bounds for $\sum'_{(\alpha, \beta)} \log(h(\alpha, \beta)^{-1})$ and for $\sum_{(\alpha, \beta)} \log M(\alpha, \beta)$, where, in the first sum, we sum over the complex solutions (α, β) of $f = g = 0$ with $h(\alpha, \beta) \neq 0$ and, in the second sum, over all complex solutions. Since

$$\begin{aligned} \sum_{(\alpha, \beta)} \log M(\alpha, \beta) &\leq \sum_{\alpha: R^{(y)}(\alpha)=0} \text{mult}(\alpha, R^{(y)}) \cdot \log(\alpha) + \sum_{\beta: R^{(x)}(\beta)=0} \text{mult}(\beta, R^{(x)}) \cdot \log(\beta) \\ &\leq \log \mathcal{M}(R^{(y)}) + \log \mathcal{M}(R^{(x)}) + 2n^2, \end{aligned}$$

it follows that $\sum_{(\alpha, \beta)} \log M(\alpha, \beta) = \tilde{O}(n^2 + n\tau)$.

For the bound on $\sum'_{(\alpha, \beta)} \log(h(\alpha, \beta)^{-1})$, we can assume that the systems $f = g = 0$ and $g = h = 0$ are in generic position such that a solution (α, β) of $f = g = 0$ is also a solution of $g = h = 0$ if and only if $\bar{R}^{(y)}(\alpha) = 0$. This can be achieved by considering a linear form $x + s \cdot y$ (and a corresponding shearing $x \mapsto x + s \cdot y$) that is separating for the union of the solutions of $f = g = 0$ and $g = h = 0$. Notice that we do not have to compute such a linear form; we only need its existence with some $s = O(\log n)$ for our argument to bound the sum $\sum'_{(\alpha, \beta)} \log(h(\alpha, \beta)^{-1})$ which is invariant with respect to the coordinate transformation.

Now write R for $R^{(y)}$ and \bar{R} for $\bar{R}^{(y)}$, and let R^* and \bar{R}^* be the corresponding square-free parts. The polynomial $q(x) := R^* / \gcd(R^*, \bar{R}^*)$ divides R^* , and the roots of q are exactly the projections of all solutions (α, β) of $f = g = 0$ for which h does not vanish. Then, for

¹¹In [39, Lem. 15], it has been shown that we can compare the roots of the resultant polynomials $\text{res}(f, f_y; y)$ and $\text{res}(f_x, f_y; y)$ with $\tilde{O}(n^6 + n^5\tau)$ bit operations. However, the proof applies to arbitrary polynomials of comparable magnitude.

each such (α, β) , it holds that

$$0 \neq \bar{R}(\alpha) = \bar{u}(\alpha, \beta) \cdot \underbrace{g(\alpha, \beta)}_{=0} + \bar{v}(\alpha, \beta) \cdot h(\alpha, \beta) = \bar{v}(\alpha, \beta) \cdot h(\alpha, \beta),$$

where \bar{u} and \bar{v} are the cofactors in the cofactor representation of \bar{R} ; see also Section 2.1. Hence, we have

$$\begin{aligned} \log(h(\alpha, \beta)^{-1}) &\leq \log(\bar{v}(\alpha, \beta)) + \log(\bar{R}(\alpha)^{-1}) \\ &\leq \tilde{O}(n\tau + n^2 \log M(\alpha, \beta) + \log(\bar{R}(\alpha)^{-1})). \end{aligned}$$

When summing over all solutions (α, β) of $f = g = 0$, the first two terms sum up to $\tilde{O}(n^4 + n^3\tau)$ since there are at most n^2 many solutions and the Mahler measure of R is bounded by $2^{O(n^2+n\tau)}$. From the definition of $q(x)$ and our genericity assumption, we conclude that the sum over the last term equals $\sum_{\alpha:q(\alpha)=0} \log(\bar{R}(\alpha)^{-1})$. Hence, since $|\bar{R}(\alpha)| \leq 2^{O(n^2+n\tau+n^2 \cdot \log(\alpha))}$ for all α , it follows that

$$\begin{aligned} \sum_{\alpha:q(\alpha)=0} \log(\bar{R}(\alpha)^{-1}) &= \sum_{\alpha:q(\alpha)=0} \log |\bar{R}(\alpha)|^{-1} + \sum_{\alpha:q(\alpha)=0} O(n^2 + n\tau + n^2 \log \alpha) \\ &= \log \left(\prod_{\alpha:q(\alpha)=0} |\bar{R}(\alpha)|^{-1} \right) + O(n^4 + n^3\tau) + n^2 \cdot \log \mathcal{M}(q) \\ &= \tilde{O}(n^4 + n^3\tau), \end{aligned}$$

where we used that q , as a divisor of R , has magnitude $(n^2, \tilde{O}(n^2 + n\tau))$ and that

$$\prod_{\alpha:q(\alpha)=0} |\bar{R}(\alpha)^{-1}| = |\text{lcf}(q(x))|^{\deg(\bar{R})} \cdot |\text{res}(q, \bar{R})|^{-1} = 2^{\tilde{O}(n^4+n^3\tau)}.$$

We conclude that the cost for evaluating the sign of h at all real valued solutions of $f = g = 0$ is bounded by $\tilde{O}(n^6 + n^5\tau)$ many bit operations.

We are left to discuss the case, where h shares a non-trivial factor with both polynomials f and g . Suppose that $p := \gcd(g, h)$ is non-trivial and define $h^* := h/p$ and $g^* := g/p$. Note that h^* and g^* are coprime, and since f and g are coprime, the same also holds for f and p . According to [39, Lem. 13], we can compute p, h^* and g^* with $\tilde{O}(n^6 + n^5\tau)$ bit operations, and the magnitude of these polynomials is bounded by $(n, O(n + \tau))$. The solutions of $f = g = 0$ now decompose into the solutions of $f = g^* = 0$ and $f = p = 0$. Trivially, $h = h^* \cdot p$ vanishes at all solutions of the latter system, hence it remains to compute the sign of h at the solutions of $f = g^* = 0$. Based on the considerations above, we can evaluate the sign of h^* as well as the sign of p at the solutions of $f = g^* = 0$ with $\tilde{O}(n^6 + n^5\tau)$ many bit operations as all involved polynomials have magnitude $(n, O(n + \tau))$.

We summarize our results:

Theorem 13. *Let $f, g \in \mathbb{Z}[x, y]$ be coprime polynomials of magnitude (n, τ) . Then, for an arbitrary polynomial $h \in \mathbb{Z}[x, y]$ of magnitude (n, τ) , we can evaluate the sign of h at all real-valued solutions of $f = g = 0$ with $\tilde{O}(n^6 + n^5\tau)$ bit operations.*

3 Computing the topology of an algebraic plane curve

Based on our results on solving bivariate polynomial systems, we will show that, using a deterministic algorithm, we can compute the topology of a planar algebraic curve

$$\mathcal{C} := \{(x, y) \in \mathbb{R}^2 : f(x, y) = 0\}$$

defined as the real zero set of a square-free integer polynomial $f \in \mathbb{Z}[x, y]$ of magnitude (n, τ) with $\tilde{O}(n^6 + n^5\tau)$ bit operations. Our algorithm can be considered as a combination of the algorithm CBISOLVE and the randomized algorithm TOPNT as introduced in [8, Sec. 3]. Hence, we only sketch our approach and refer to [8] for more details. Since the complexity analysis of our algorithm crucially depends on the results presented in [32, 39], we suggest to also consult these papers for amortized bounds on the complexity of the considered computations. As input, the algorithm receives the exact integer coefficients of the polynomial f . Considering the following consecutive computations, it eventually returns a planar straight-line graph \mathcal{G} (embedded in \mathbb{R}^2) that is isotopic to \mathcal{C} :

1. Compute $f_x^* := f_x / \gcd(f_x, f_y)$ and $f_y^* := f_y / \gcd(f_x, f_y)$.
2. Determine an integer s of absolute value less than n^4 such that
 - (a) $l_s = x + s \cdot y$ is a separating form for the *strongly critical* points of f , that is, the solutions of the system $f_x^* = f_y^* = 0$, and such that
 - (b) the leading coefficient of $f(x + sy, y)$ with respect to y is a constant.

We use the term *strongly critical* to denote the additional restriction $f_x^* = f_y^* = 0$ over *critical* points, which satisfy $f_x = f_y = 0$.

3. Perform the coordinate transformation (*shearing*) $x \mapsto x + s \cdot y$, that is, replace the polynomial $f(x, y)$ by $F(x, y) := f(x + sy, y)$. We define

$$\bar{\mathcal{C}} := \{(x, y) \in \mathbb{R}^2 : F(x, y) = 0\}$$

to be the real vanishing set of the polynomial $F(x, y)$. Note that the shearing does not change the isotopy of the curve; hence, it will suffice to compute a straight-line graph isotopic to $\bar{\mathcal{C}}$.

4. Isolate all real valued solutions (α_i, β_i) , with $i = 1, \dots, k$ and some $k \leq n^2$, of the polynomial system $F = F_y = 0$. The points (α_i, β_i) and the (not necessarily distinct) values α_i are called *x-critical points* and *x-critical values* of $\bar{\mathcal{C}}$, respectively. W.l.o.g., we can assume that $\alpha_1 \leq \dots \leq \alpha_k$.
5. Compute the sign of F_x at all *x-critical points* $(\alpha_i, \beta_i) \in \mathbb{R}^2$ of $\bar{\mathcal{C}}$. Each *x-critical point* (α_i, β_i) with $F_x(\alpha_i, \beta_i) = 0$ is a *singular point* of $\bar{\mathcal{C}}$.
6. Isolate all real roots γ_j of the polynomial

$$\hat{R}(x) := \frac{\partial R^*}{\partial x} \Big/ \gcd\left(\frac{\partial R^*}{\partial x}, \frac{\partial^2 R^*}{\partial x^2}\right)$$

which is the square-free part of the derivative of R^* , and $R^* := R / \gcd(R, \frac{\partial R}{\partial x})$ is defined as the square-free part of the resultant $R := \text{res}(F, F_y; y)$.

The values $\gamma_1, \dots, \gamma_m$ of \hat{R} separate the roots of R , that is, in between of two distinct consecutive roots α_i and α_{i+1} of R , there exists at least one root γ_j of \hat{R} . If necessary, we discard some γ_j in an arbitrary manner such that, in between of two consecutive distinct values α_i and α_{i+1} , there exists exactly one γ_i . (Notice that it is not guaranteed that the \bar{C} is in generic position with respect to its x -critical values, that is, there might be two distinct x -critical points (α_i, β_i) and $(\alpha_{i+1}, \beta_{i+1})$ with $\alpha_i = \alpha_{i+1}$.)

We remark that, for a practical implementation of the algorithm, we propose to consider arbitrary rational values γ_i in between each pair of consecutive real roots of R . However, in general, this yields a worse complexity bound for the root isolation in the next step.

7. Isolate the real roots of all polynomials $F(\alpha_i, y)$ and $F(\gamma_i, y)$.
8. Connect points $(\alpha_i, y^*) \in \bar{C}$ (or $(\alpha_{i+1}, y^*) \in \bar{C}$) and $(\gamma_i, y^{**}) \in \bar{C}$ by a line segment if and only if they are connected via an arc of \bar{C} . Return the so obtained planar straight-line graph \mathcal{G} , which is isotopic to \bar{C} , and thus also isotopic to \mathcal{C} .

In the previous sections, we have shown how to perform the first five steps (except for Step 2 (b)) with $\tilde{O}(n^6 + n^5\tau)$ bit operations, considering that the polynomials f, f_x^* and f_y^* as well as the transformed polynomials $F(x, y) = f(x + sy, y)$, $F_x(x, y) = f_x(x + sy, y)$ and $F_y(x, y) = s \cdot f_x(x + sy, y) + f_y(x, y)$ have magnitude $(n, O(n \log n + \tau))$, because the bitsize of s is bounded by $O(\log n)$. For the computation needed to guarantee 2 (b), we remark that all except n bad values for s fulfill 2 (b) and that we can compute these values with a number of bit operations bounded by $\tilde{O}(n^3 + n^2\tau)$. Namely, these bad values are exactly the roots of the leading coefficient $f_n^{(y)}(s) \in \mathbb{Z}[s]$ of $f(x + sy, y) \in \mathbb{Z}[s, x][y]$ with respect to y . Hence, the computation of $f_n^{(y)}(s)$ as well as approximating all of its roots to an error of less than $1/2$ needs no more than $\tilde{O}(n^3 + n^2\tau)$ bit operations. For Step 6, we refer to [39, Sec. 3], where a bit complexity bound of $\tilde{O}(n^6 + n^5\tau)$ has been given.

The last two steps of the above algorithm need more explanation: According to [8, Sec. 3.2.2], there exists no solution of $f = \gcd(f_x, f_y) = 0$ in \mathbb{C}^2 . That is, there are either no solutions (iff $\gcd(f_x, f_y)$ is trivial) or all solutions are located at infinity. Hence, all singular points of \mathcal{C} are the common solutions of $f = f_x^* = f_y^* = 0$. From the transformation in Step 3, we conclude that there are no two complex solutions of $F_x^* = F_y^* = 0$ (i.e. strongly critical points of F) sharing the same x -coordinate, where $F_x^* := F_x / \gcd(F_x, F_y)$ and $F_y^* := F_y / \gcd(F_x, F_y)$. Namely, the *strongly* critical points of F are directly obtained from shearing the strongly critical points of f . Notice that this does not hold for the x -critical points of f and F ; in particular, choosing a separating form $x + sy$ for the system $f = f_y = 0$ does not imply that the sheared curve \bar{C} is in generic position with respect to its x -critical points. However, if we choose s according to the requirements in Step 2, then \bar{C} is in generic position with respect to its *strongly* critical points. In particular, this means that for each x -critical value α_i of \bar{C} , there exists at most one strongly critical point above α_i . Furthermore, from the computation in Step 5, we can also determine whether there exists a singular point with x -coordinate α_i or not. At this point, we remark that each singular point above α_i must be real valued because there exists at most one strongly critical point above α_i and F has real

valued coefficients, which implies that the critical points above α must arise in complex conjugate pairs. We can then use the computation in [8, Sec. 3.2.2.] to determine the number n_{α_i} of distinct complex roots of $F(\alpha_i, y)$. Namely, it holds that

$$n_{\alpha_i} = \begin{cases} n - \text{mult}(\alpha_i, R) + \text{mult}(\alpha_i, Q), & \text{if there exists a singular point } (\alpha_i, \beta) \in \bar{\mathcal{C}} \\ n - \text{mult}(\alpha_i, R), & \text{if there exists no singular point } (\alpha_i, \beta) \in \bar{\mathcal{C}}, \end{cases}$$

where $R := \text{res}(F, F_y; y)$ and $Q := \text{res}(F_x^*, F_y^*; y)$. In other words, n_{α_i} can be derived from the multiplicity of α_i as a root of R (and Q) and the fact whether there exists a singular point above α_i or not. For computing the values $\text{mult}(\alpha_i, R)$ and $\text{mult}(\alpha_i, Q)$ for all α_i , we refer to [39, Lem. 15], where it has been shown that the cost for the necessary computations is bounded by $\tilde{O}(n^6 + n^5\tau)$ bit operations. Since each polynomial $F(\gamma_i, y)$ has degree n and all its roots are simple, the number of distinct roots of $F(\gamma_i, y)$ equals n for all i . Now, having computed the number of distinct roots for each of the polynomials $F(\alpha_i, y)$ and $F(\gamma_i, y)$, we can use the algorithm from [39] to compute isolating intervals for all real roots of the latter polynomials together with the corresponding multiplicities. Using Theorem 7, one can show that the cost for the root isolation and for computing sufficiently good approximations of the roots of R and \hat{R} is bounded by $\tilde{O}(n^6 + n^5\tau)$ bit operations; for more details, we refer to [39, Lem. 20]. After Step 7, we have already computed the vertices of the planar graph \mathcal{G} , and each of these vertices is located on the curve $\bar{\mathcal{C}}$. It remains to show how to connect the vertices by line segments in an appropriate manner. That is, we have to determine whether two vertices in neighboring fibers are connected via an arc of $\bar{\mathcal{C}}$ and then connect them by a line segment if and only if the latter is the case. Notice that, for each value $x = \gamma_i$ that is not x -critical, we have computed isolating intervals for all real roots of the polynomial $F(\gamma_i, y)$, and each such root is simple. For each x -critical value $x = \alpha_i$, we have also computed isolating intervals for all real roots $y_{i,1}, \dots, y_{i,m_i}$ of $F(\alpha_i, y)$, with $m_i \leq n_{\alpha_i}$, and, in addition, we know

- the multiplicity $\mu_{i,j} := \text{mult}(y_{i,j}, F(\alpha_i, y))$ of each root $y_{i,j}$,
- the sign of $F_x(\alpha_i, y_{i,j})$ for each $y_{i,j}$ with $\mu_{i,j} > 1$, and, in particular,
- whether $(\alpha_i, y_{i,j})$ is a singular point or not.

For each x -critical point $(\alpha_i, y_{i,j})$ that is not singular, we can further compute the sign of the $\mu_{i,j}$ -th partial derivative $F_y^{(\mu_{i,j})}$ at $(\alpha_i, y_{i,j})$ with respect to y by approximating $(\alpha_i, y_{i,j})$ and approximately evaluating $F_y^{(\mu_{i,j})}(\alpha_i, y_{i,j})$, where we define $F_y^{(k)} := \frac{\partial^k}{\partial y^k} F$. The analysis in [39] shows that the necessary computations can be carried out using $\tilde{O}(n^6 + n^5\tau)$ bit operations. More precisely, the cost for evaluating F at some point (x_0, y_0) with an output precision of L bits after the binary point is bounded by $\tilde{O}(n^2(L + n + \tau + n \log \max(|x_0|, |y_0|)))$ and the sum of all values $\log(F_y^{(\mu_{i,j})}(\alpha_i, y_{i,j})^{-1})$ and $\log \max(|\alpha_i|, |y_{i,j}|)$ is bounded by $\tilde{O}(n^4 + n^3\tau)$; see [39, Lem. 17 and 18]. Also, the cost for computing good enough approximations of the roots $y_{i,j}$ is bounded by $\tilde{O}(n^6 + n^5\tau)$; see the proof of [39, Lem. 20] for details.

We claim that the above information is already sufficient to determine (in a purely combinatorial way) whether two points in the fibers $x = \alpha_i$ and $x = \gamma_i$ are connected via an arc of $\bar{\mathcal{C}}$: Namely, the local topology of $\bar{\mathcal{C}}$ at each point that is not x -critical is trivial, that is, there exists

exactly one arc that enters the point from the left and leaves to the right. For each x -critical point $p = (\alpha_i, \beta_i)$ that is not singular, there is exactly one arc incident to p as well, and one of the following three possibilities for the local topology at p ; see also Figure 1 for an illustration:

- (Case 1) The arc enters p from the left and leaves to the right,
- (Case 2) the arc enters p from the left and leaves to the left, or
- (Case 3) the arc enters p from the right and leaves to the right.

In the latter two cases, the point p is an x -extremal point, that is, the curve $\tilde{\mathcal{C}}$ makes a turn at p . In the first case, the local topology of $\tilde{\mathcal{C}}$ is the same as at any point that is not x -critical, however, the arc passing through p is vertical at p . Obviously, the first case applies if and only if the multiplicity $\mu = \text{mult}(\beta_i, F(\alpha_i, y))$ of β_i as a root of $F(\alpha_i, y)$ is odd. If μ is even, then case 2 applies if and only if $F_x(\alpha_i, \beta_i) \cdot F_y^{(\mu)}(\alpha_i, \beta_i) > 0$, and case 3 applies if and only if $F_x(\alpha_i, \beta_i) \cdot F_y^{(\mu)}(\alpha_i, \beta_i) < 0$.

In summary, for each point $p = (\alpha_i, \beta_i)$ in the fiber $x = \alpha_i$ except for a unique singular point (if such a point exists), we know the local topology of $\tilde{\mathcal{C}}$ at p . Hence, we can determine all connections between the non-singular points in an x -critical fiber $x = \alpha_i$ and the points in a neighboring fiber $x = \gamma_i$ (that is not x -critical) from bottom to top and from top to bottom until only the connections to the singular point are left. Then, all points in the fiber $x = \gamma_i$ that are not connected yet with a point in the fiber $x = \alpha_i$ must be connected with the unique singular point.

We conclude that the topology of \mathcal{C} can be computed with a number of bit operations bounded by $\tilde{O}(n^6 + n^5\tau)$. Since we can always compute the square-free part f^* of an arbitrary bivariate polynomial $f \in \mathbb{Z}[x, y]$ of magnitude (n, τ) with $\tilde{O}(n^6 + n^5\tau)$ many bit operations without affecting the zero set, and since f^* has magnitude $(n, O(n \log n + \tau))$, we obtain the following general result:

Theorem 14. *Given an arbitrary (not necessarily square-free) polynomial $f \in \mathbb{Z}[x, y]$ of magnitude (n, τ) , we can compute the topology of the real planar algebraic curve*

$$\mathcal{C} := \{(x, y) \in \mathbb{R}^2 : f(x, y) = 0\}$$

with a number of bit operations bounded by $\tilde{O}(n^6 + n^5\tau)$.




Figure 1: Connecting points in a critical fiber $x = \alpha_i$ (on the left) with points in a neighboring non-critical fiber $x = \gamma_i$ (on the right) is purely combinatorial if there exists at most one singular point (blue diamond) and the local topology at the non-singular x -critical points (red squares) is known.

A Missing Proofs

Theorem 9. Let $F \in \mathbb{Z}[x]$ be an integer polynomial of magnitude (d, μ) with distinct roots $z_1, \dots, z_m \in \mathbb{C}$. Then, it holds that

$$\sum_{i=1}^m \text{mult}_i \log \text{sep}(z_i, F)^{-1} = \tilde{O}(d^2 + d\mu).$$

In particular,

$$\Sigma^*(F) \leq \Sigma(F) = \tilde{O}(d^2 + d\mu).$$

Proof. We consider the factorization of F (over \mathbb{Z}) into square-free and pair-wise coprime factors:

$$F(x) = \prod_{i=1}^k Q_i(x)^{s_i}, \text{ with } d_i := \deg(Q_i) \geq 1,$$

such that the polynomials $Q_i(x)$ and $F(x)/Q_i(x)^{s_i}$ are coprime, and $d = \sum_{i=1}^k d_i s_i$. We further denote F^* the square-free part of F and $m := \deg(F^*) = \sum_{i=1}^k d_i$ its degree. Then, for arbitrary roots α and β of F^* , it holds that

$$\begin{aligned} |(F^*)'(\alpha)| &= |\text{lcf}(F^*)| \cdot |\alpha - \beta| \cdot \prod_{\gamma \neq \alpha, \beta: F^*(\gamma)=0} |\gamma - \alpha| \\ &\leq |\text{lcf}(F^*)| \cdot |\alpha - \beta| \cdot \prod_{\gamma \neq \alpha, \beta: F^*(\gamma)=0} 2M(\alpha, \gamma) \\ &\leq 2^{m-2} \cdot |\alpha - \beta| \cdot M(\alpha)^{m-3} \cdot \mathcal{M}(F^*) \end{aligned}$$

since $\mathcal{M}(F^*) = |\text{lcf}(F^*)| \cdot \prod_{z: F^*(z)=0} M(z)$. Suppose, w.l.o.g., that α is a root of Q_i and β is a root of F^* closest to α . Then, according to the above inequality, we have

$$\text{sep}(\alpha, F) = |\alpha - \beta| \geq \frac{|(F^*)'(\alpha)|}{2^{m-2} \cdot M(\alpha)^{m-3} \cdot \mathcal{M}(F^*)}.$$

We now apply this inequality to the product over all $\text{sep}(\alpha_j, F)$, $j = 1, \dots, d_i$, where $\alpha_1, \dots, \alpha_{d_i}$ denote the roots of Q_i :

$$\begin{aligned} \prod_{j=1}^{d_i} \text{sep}(\alpha_j, F) &\geq 2^{(2-m)d_i} \cdot \mathcal{M}(Q_i)^{3-m} \cdot \mathcal{M}(F^*)^{-d_i} \cdot \prod_{j=1}^{d_i} |(F^*)'(\alpha_j)| \\ &= 2^{(2-m)d_i} \cdot \mathcal{M}(Q_i)^{3-m} \cdot \mathcal{M}(F^*)^{-d_i} \cdot \prod_{j=1}^{d_i} \left| (Q_i)'(\alpha_j) \cdot \frac{F^*}{Q_i}(\alpha_j) \right| \end{aligned} \tag{18}$$

since

$$(F^*)'(\alpha_j) = \underbrace{Q_i(\alpha_j)}_{=0} \cdot \left(\frac{F^*}{Q_i} \right)'(\alpha_j) + (Q_i)'(\alpha_j) \cdot \frac{F^*}{Q_i}(\alpha_j).$$

In addition, we have

$$\prod_{j=1}^{d_i} |Q'_i(\alpha_j)| = |\text{lcf}(Q_i)^{1-d_i} \cdot \text{res}(Q_i, Q'_i)| \geq |\text{lcf}(Q_i)|^{1-d_i} \quad \text{and}$$

$$\prod_{j=1}^{d_i} \left| \frac{F^*}{Q_i}(\alpha_j) \right| = \left| \text{lcf}(Q_i)^{d_i-m} \cdot \text{res}\left(Q_i, \frac{F^*}{Q_i}\right) \right| \geq |\text{lcf}(Q_i)|^{d_i-m},$$

since $\text{res}(Q_i, Q'_i)$ and $\text{res}(Q_i, \frac{F^*}{Q_i})$ are non-zero integers. Applying the latter two inequalities to (18) now yields

$$\prod_{j=1}^{d_i} \text{sep}(\alpha_j, F) \geq 2^{(2-m)d_i} \cdot \mathcal{M}(Q_i)^{3-d} \cdot \mathcal{M}(F^*)^{-d_i} \cdot |\text{lcf}(Q_i)|^{1-m}.$$

Finally, consider the product of the separations of all roots to the respective powers s_i :

$$\begin{aligned} \prod_{i=1}^k \prod_{j=1}^{d_i} \text{sep}(\alpha_j, F)^{s_i} &\geq \prod_{i=1}^k 2^{(1-m)d_i s_i} \cdot \mathcal{M}(Q_i)^{(3-m)s_i} \cdot \mathcal{M}(F^*)^{-d_i s_i} \cdot \prod_{i=1}^k |\text{lcf}(Q_i)|^{-s_i} \\ &= 2^{(1-m)d} \cdot \mathcal{M}(F)^{3-m} \cdot \mathcal{M}(F^*)^{-d} \cdot |\text{lcf}(F)|^{-1} = 2^{-\tilde{O}(d^2 + d\mu)}, \end{aligned}$$

where we used that $\prod_{i=1}^k \mathcal{M}(Q_i)^{s_i} = \mathcal{M}(F)$ by the multiplicativity of the Mahler measure and $\mathcal{M}(F^*) \leq \mathcal{M}(F) = 2^{O(\mu + \log d)}$. This shows that $\sum_{j=1}^m \log \text{sep}(z_j, F)^{-1} = \tilde{O}(d^2 + d\mu)$. Since, for each root z_j of F , $\text{sep}(z_j, F)$ is upper bounded by two times the maximal absolute value of the roots of F , we have $\text{sep}(z_j, F) < 2^{\mu+2}$ according to the Cauchy root bound. Thus, it follows that $\sum_{j=1}^m \log \text{sep}(z_j, F)^{-1} \leq d(\mu + 2) + \sum_{j=1}^m \log \text{sep}(z_j, F)^{-1} = \tilde{O}(d^2 + d\mu)$. \square

Lemma 15. *Suppose that the properties (Isol 1) and (Isol 2) are fulfilled. Then, the values $\text{LB}(\alpha)$ and $\text{LB}(\beta)$ as defined in (10) constitute lower bounds for the absolute values of the resultants $R^{(y)}$ and $R^{(x)}$ restricted to the boundary of $D(\alpha)$ and $D(\beta)$, respectively.*

Proof. The proof is almost identical to the proof of [6, Lem. 3.1]. Write $m := m(\alpha)$, $r := r(\alpha)$, and $R := R^{(y)}$. According to (Isol 1), $D_{r/2}(m)$ is also isolating, hence $\frac{|z-\alpha|}{|(m-r)-\alpha|} > \frac{r/2}{3r/2} > \frac{1}{4}$ for all points $z \in \partial D(\alpha)$ on the boundary of $D(\alpha)$. In addition, (Isol 2) guarantees that $R(m-r) \neq 0$ and that, for any root $\alpha' \neq \alpha$ of R , it holds that $\frac{|z-\alpha'|}{|(m-r)-\alpha'|} \geq \frac{|\alpha'-m|-r}{|\alpha'-m|+r} \geq 1 - \frac{2r}{4r} = \frac{1}{2}$. Hence, it follows that

$$\begin{aligned} \frac{|R(z)|}{|R(m-r)|} &= \left(\frac{|z-\alpha|}{|(m-r)-\alpha|} \right)^{\text{mult}_\alpha} \cdot \prod_{\zeta \neq \alpha: R(\zeta)=0} \left(\frac{|z-\alpha'|}{|(m-r)-\alpha'|} \right)^{\text{mult}'_{\alpha'}} \\ &> 4^{-\text{mult}_\alpha} 2^{-(\deg R - \text{mult}_\alpha)} = 2^{-\text{mult}_\alpha - \deg R}. \end{aligned} \quad \square$$

Lemma 16. Let $\text{LB}(\alpha)$ be defined as in (10). Then,

$$\sum_{\alpha} \overline{\log}(\text{LB}(\alpha)^{-1}) = \tilde{O}(N^2 + NT),$$

where we sum over all distinct complex roots of the resultant polynomial $R^{(y)}$.

Proof. For any root α of the resultant $R := R^{(y)}$, write $m := m(\alpha)$, $r := r(\alpha)$, and

$$R(x) = \text{lcf } R \cdot (x - \alpha)^{\text{mult}_{\alpha}} \cdot \prod_{\zeta \neq \alpha} (x - \zeta)^{\text{mult}_{\zeta}},$$

where ζ runs over all distinct complex roots of R but α . Define $\sigma_{\alpha} := \min\{\text{sep}_{\alpha}, M(\alpha)\}$. Property (Isol 3) ensures that $\sigma_{\alpha}/16 \leq r \leq \sigma_{\alpha}/4$. Hence, it follows that

$$\begin{aligned} \frac{1}{32}\sigma_{\alpha} &\leq \frac{1}{2}r \leq |(m-r) - \alpha| \leq \frac{3}{2}r \leq \frac{3}{8}\sigma_{\alpha} < \frac{1}{2}\sigma_{\alpha} \quad \text{and} \\ \frac{1}{2}|\alpha - \zeta| &< |\alpha - \zeta| - \frac{1}{2}r \leq |(m-r) - \zeta| \leq |\alpha - \zeta| + \frac{3}{2}r \leq 2|\alpha - \zeta| \end{aligned}$$

for all roots $\zeta \neq \alpha$ of R . Hence,

$$\begin{aligned} \text{LB}(\alpha) &\leq 2^{-\text{mult}_{\alpha} - \deg R} \cdot |R(m-r)| \\ &< 2^{-N} \cdot |\text{lcf } R| \cdot |(m-r) - \alpha|^{\text{mult}_{\alpha}} \cdot \prod_{\zeta \neq \alpha} |(m-r) - \zeta|^{\text{mult}_{\zeta}} \\ &< 2^{-N} \cdot (\frac{1}{2}\sigma_{\alpha})^{\text{mult}_{\alpha}} \cdot |\text{lcf } R| \cdot \prod_{\zeta \neq \alpha} (2|\alpha - \zeta|)^{\text{mult}_{\zeta}} \\ &< \sigma_{\alpha}^{\text{mult}_{\alpha}} \cdot |\text{lcf } R| \cdot \prod_{\zeta \neq \alpha} |\alpha - \zeta|^{\text{mult}_{\zeta}} \\ &= \sigma_{\alpha}^{\text{mult}_{\alpha}} \cdot \frac{|R^{(\text{mult}_{\alpha})}(\alpha)|}{\text{mult}_{\alpha}!} \\ &= \sigma_{\alpha}^{\text{mult}_{\alpha}} \cdot 2^{O(N+T)} \cdot M(\alpha)^N \\ &= 2^{O(\text{mult}_{\alpha}(N+T))} M(\alpha)^N, \end{aligned} \tag{19}$$

since $R^{(\text{mult}_{\alpha})}/\text{mult}_{\alpha}! \in \mathbb{Z}[x]$ has magnitude $(N, O(N+T))$ and $\sigma_{\alpha} \leq \max_{\zeta: R(\zeta)=0} |\zeta| = 2^{O(N+T)}$ according to Cauchy's root bound, where the maximum is taken over all roots of R including α .

We can also compute a lower bound for $\text{LB}(\alpha)$:

$$\begin{aligned} \text{LB}(\alpha) &\geq \frac{1}{2} \cdot 2^{-\text{mult}_{\alpha} - \deg R} \cdot |R(m-r)| \\ &\geq 2^{-2N-1} \cdot |\text{lcf } R| \cdot (\frac{1}{32}\sigma_{\alpha})^{\text{mult}_{\alpha}} \cdot \prod_{\zeta \neq \alpha} (\frac{1}{2}|\alpha - \zeta|)^{\text{mult}_{\zeta}} \\ &> 2^{-8N} \cdot \sigma_{\alpha}^{\text{mult}_{\alpha}} \cdot |\text{lcf } R| \cdot \prod_{\zeta \neq \alpha} |\alpha - \zeta|^{\text{mult}_{\zeta}}. \end{aligned} \tag{20}$$

Since we are mainly interested in a bound for the product over all $\text{LB}(\alpha)$, we first consider the product

$$\Pi := \prod_{\alpha} (2^{-8N} \cdot \sigma_{\alpha}^{\text{mult}_{\alpha}} \cdot |\text{lcf } R| \cdot \prod_{\zeta \neq \alpha} |\alpha - \zeta|^{\text{mult}_{\zeta}})$$

of the bound in (20) over all α . Since $\sum_{\alpha} \text{mult}_{\alpha} \leq N$, it follows that $\prod_{\alpha} 2^{-8N} = 2^{O(N^2)}$. For the product of the remaining factors, we first write the square-free decomposition of

R as $R = \prod_s r_s^s$ with square-free, pairwise coprime $r_s \in \mathbb{Z}[x]$. Since $R^{(s)}/s!$ has integer coefficients, we have

$$1 \leq \left| \operatorname{res} \left(r_s, \frac{R^{(s)}}{s!} \right) \right| = |\operatorname{lcf} r_s|^{\deg R - s} \cdot \prod_{\zeta: r_s(\zeta)=0} \frac{R^{(s)}(\zeta)}{s!}$$

and, thus,

$$\begin{aligned} & \prod_{\alpha} (\sigma_{\alpha}^{\operatorname{mult}_{\alpha}} \cdot |\operatorname{lcf} R| \cdot \prod_{\zeta \neq \alpha} |\alpha - \zeta|^{\operatorname{mult}_{\zeta}}) \\ & > 2^{-\Sigma(R)} \cdot \prod_{\alpha} |\operatorname{lcf} R| \prod_{\zeta \neq \alpha} |\alpha - \zeta|^{\operatorname{mult}_{\zeta}} \\ & = 2^{-\Sigma(R)} \cdot \prod_{\alpha} \frac{|R^{(\operatorname{mult}_{\alpha})}(\alpha)|}{\operatorname{mult}_{\alpha}!} \\ & = 2^{-\Sigma(R)} \cdot \prod_s |\operatorname{lcf} r_s|^{s-N} \cdot \left| \operatorname{res} \left(r_s, \frac{R^{(s)}}{s!} \right) \right| \\ & \geq 2^{-\Sigma(R)} \cdot |\operatorname{lcf} R| \cdot |\operatorname{lcf} R^*|^{-N} \\ & = 2^{-\tilde{O}(N^2+NT)}, \end{aligned}$$

where we used that $\Sigma(R) = \tilde{O}(N^2 + NT)$. Hence, Π is lower bounded by $2^{-\tilde{O}(N^2+NT)}$. Similar to the computation in (19), we can also determine an upper bound for the factor in Π corresponding to an arbitrary but fixed α . Namely, $\sigma_{\alpha}^{\operatorname{mult}_{\alpha}} = 2^{O(\operatorname{mult}_{\alpha}(N+T))}$ and

$$|\operatorname{lcf} R| \prod_{\zeta \neq \alpha} |\alpha - \zeta|^{\operatorname{mult}_{\zeta}} = \frac{|R^{(\operatorname{mult}_{\alpha})}(\alpha)|}{\operatorname{mult}_{\alpha}!} = 2^{O(T+\log N)} M(\alpha)^N.$$

Thus, for any subset A of distinct roots of R , the partial product

$$\Pi' := \prod_{\alpha \in A} (2^{-8N} \cdot \sigma_{\alpha}^{\operatorname{mult}_{\alpha}} \cdot |\operatorname{lcf} R| \cdot \prod_{\zeta \neq \alpha} |\alpha - \zeta|^{\operatorname{mult}_{\zeta}})$$

is upper bounded by $2^{\tilde{O}(N^2+NT)} \prod_{\alpha \in A} M(\alpha)^N = 2^{\tilde{O}(N^2+NT)}$ since $\prod_{\alpha \in A} M(\alpha) = 2^{O(T)}$; cf. Lemma 8. We conclude that

$$\sum_{\alpha} \log(\operatorname{LB}(\alpha)^{-1}) \leq \tilde{O}(N^2 + NT) + \sum_{\alpha} \log(\operatorname{LB}(\alpha)^{-1}). \quad \square$$

B Fast Approximate Polynomial Multipoint Evaluation

Given a non-negative integer $L \in \mathbb{N}$, a polynomial $F(x) = \sum_{i=0}^n F_i x^i \in \mathbb{C}[x]$ of degree n and complex points $x_1, \dots, x_n \in \mathbb{C}$, the task of approximate polynomial multipoint evaluation is to compute approximations \tilde{y}_j for $y_j := F(x_j)$ such that $|\tilde{y}_j - y_j| \leq 2^{-L}$ for all $j = 1, \dots, n$. Let 2^τ and 2^Γ , with $\tau, \Gamma \in \mathbb{N}_{\geq 1}$, denote bounds on the absolute values of the coefficients of F and the points x_j , respectively.

We aim to show that, using approximate arithmetic in the classical fast polynomial multipoint evaluation algorithm from [40] (see also [27, Sec. 10.1]), we can compute approximations \tilde{y}_i as above with $\tilde{O}(n(L + \tau + n\Gamma))$ bit operations, and it suffices to consider $L + \tilde{O}(\tau + n\Gamma)$ bits of the coefficients of F and the points x_i . In particular, if L dominates τ and $n\Gamma$, the precision demand is essentially linear in L , and the computation time is linear in L and n .

This fact has been observed previously by Kirrinnis [33, Thm. 3.9 and App. A.3] in a slightly different context. Unfortunately, the formulation therein is less general and requires some transformations of the input, and the result does not seem to be widely known in the community. Both Kirrinnis' and our discussion in [35] rely on a fast numerical polynomial division scheme using fast Fourier transforms due to Schönhage [47]. Existing implementations suggest that this algorithm is only efficient for extraordinarily large inputs. Thus, we describe a different analysis based on polynomial division via Newton's method for polynomial inversion, which is known to have a more moderate break-even point. To the best of our knowledge, this scheme has not been extensively analyzed in a numerical setting before, and we believe this section to be of independent interest.

For the sake of simplicity, assume that $n = 2^k$ is a power of two; otherwise, pad F with zeros. We require that arbitrarily good approximations of the coefficients F_i and the points x_j are provided by an oracle for the cost of reading the approximations. That is, asking for an approximation of the coefficients of F and all points x_j to a precision of ℓ bits after the binary point takes $O(n(\tau + \ell))$ and $O(n(\Gamma + \ell))$ bit operations, respectively.

Algorithm 17 (Multipoint evaluation). We will follow the classical divide-and-conquer method for fast polynomial multipoint evaluation [40, 27]:

1. From the linear factors $g_{0,j}(x) := x - x_j$, we recursively compute the subproduct tree

$$g_{i,j}(x) := (x - x_{(j-1)2^i+1}) \cdots (x - x_{j2^i}) = g_{i-1,2j-1}(x) \cdot g_{i-1,2j}(x) \quad (21)$$

for i from 1 to $k-1$ and j from 1 to $n/2^i = 2^{k-i}$, that is, going up from the leaves. Notice that $\deg g_{i,j} = 2^i$.

2. Starting with $r_{k,1}(x) := F(x)$, we recursively compute the remainder tree

$$r_{i,j}(x) := F(x) \bmod g_{i,j}(x) = r_{i+1,\lceil j/2 \rceil}(x) \bmod g_{i,j}(x)$$

for i from $k-1$ to 0 and j from 1 to $n/2^i = 2^{k-i}$, that is, going down from the root. Notice that $\deg r_{i,j} < 2^i$.

3. Observe that the value at point x_j is exactly the remainder

$$r_{0,j} = F(x) \bmod g_{0,j}(x) = F(x) \bmod (x - x_j) = F(x_j) \in \mathbb{C}.$$

For the polynomial division with remainder, we use an asymptotically fast recursive approach, often called Newton's method for polynomial inversion. It relies on Hensel lifting to compute the inverse of the *reverse* polynomial (see page 7 for the definition) of the divisor modulo some power of x , which translates to the quotient in the original division.

Algorithm 18. Given a polynomial $F = \sum_{i=0}^{2n} F_i x^i \in \mathbb{C}[x]$ of degree at most $2n$ and a *monic* polynomial $G = \sum_{i=0}^n G_i x^i \in \mathbb{C}[x]$ of degree n , we compute the quotient Q and the remainder R of the polynomial division of F by G in the following way:

1. Define $f := \text{rev}_{2n} F$, $g := \text{rev}_n G$, $h_0 := 1$, and $k := \lceil \log(n+1) \rceil$.
2. For $i = 1, \dots, k$, recursively compute $h_i := 2h_{i-1} - g \cdot h_{i-1}^2 \bmod x^{2^i}$.
3. Compute $q := f \cdot h_k \bmod x^{n+1}$, $Q = \text{rev}_n q$, and $R = F - Q \cdot G$.

The correctness of the above algorithm follows from the loop invariant that h_i is a multiplicative inverse of g modulo x^{2^i} : Observe that $h_0 \cdot g \equiv 1 \bmod x$ since g has constant coefficient 1. By definition, $h_i \cdot g \equiv 1 - (h_{i-1} \cdot g - 1)^2 \bmod x^{2^i}$, and thus, by induction, we have $h_{i-1} \cdot g - 1 \equiv 0 \bmod x^{2^{i-1}}$. It follows that $h_i \cdot g \equiv 1 \bmod x^{2^i}$. It is now straightforward to verify that R as defined in Step 3 is of degree at most $n-1$, and thus, Q and R are indeed the unique quotient and remainder of the division of F by G .

We further remark that $h_i \equiv h_{i-1} \bmod x^{2^{i-1}}$ for all $i > 0$.

The arithmetic complexity of Algorithm 18, counting exact additions and multiplications in \mathbb{C} , is $O(\text{Mul}(n))$, where $\text{Mul}(n)$ denotes the arithmetic complexity of multiplication of two n -th degree polynomials. For Algorithm 17, it follows that, in the i -th layer of the subproduct tree and the remainder tree, a total number of $2^{k-i} \cdot O(\text{Mul}(2^i)) = O(\text{Mul}(n))$ field operations suffices. Thus, the arithmetic complexity of multipoint evaluation is $O(\text{Mul}(n) \cdot \log n)$, which simplifies to $\tilde{O}(n)$ if quasi-linear time polynomial multiplication algorithms are used [48, 26, 18].

To derive bounds on the bit complexity of these methods, when applied with approximate arithmetic, we recall the bit complexity and precision demand of approximate polynomial multiplication. In the following considerations, we stipulate that \tilde{P} is an L -(bit) approximation of some P if $\|\tilde{P} - P\|_1 \leq 2^{-L}$.

Lemma 19. Let F and $G \in \mathbb{C}[x]$ be polynomials of magnitude bounded by (n, τ) , where $\tau \in \mathbb{N}$. Computing an ℓ -bit approximation of $H = F \cdot G$, that is, computing an $\tilde{H} \in \mathbb{C}[x]$ such that $\|\tilde{H} - H\|_1 \leq 2^{-\ell}$, is possible in

$$O(\text{Mul}(n(\ell + \tau + 2 \log n))) \quad \text{or} \quad \tilde{O}(n(\ell + \tau))$$

bit operations and with a precision demand of at most $\ell + O(\tau + \log n)$ bits on each of the coefficients of F and G .

Proof. Let $s := \ell + \tau + 2\lceil \log(n+1) \rceil + 2$. Define $f := 2^s F$ and $g := 2^s G$, and notice that $h := f \cdot g = 2^{2s} H$. We consider polynomials \tilde{f} and $\tilde{g} \in \mathbb{Z}[i][x]$ obtained from f and g by truncating the coefficients after the binary point, and write $\Delta f := \tilde{f} - f$ and $\Delta g := \tilde{g} - g$. Since $\|\Delta f\|_1, \|\Delta g\|_1 \leq n+1$ by definition of \tilde{f} and \tilde{g} ,

$$\begin{aligned} \|\tilde{f}\tilde{g} - fg\|_1 &\leq \|\Delta f\|_1 \cdot \|g\|_1 + \|f\|_1 \cdot \|\Delta g\|_1 + \|\Delta f\|_1 \cdot \|\Delta g\|_1 \\ &\leq (n+1)^2 \cdot 2^{s+\tau} + (n+1)^2 \cdot 2^{s+\tau} + (n+1)^2 \leq (n+1)^2 \cdot 2^{s+\tau+2} \end{aligned}$$

holds. For $\tilde{H} := 2^{-2s} \tilde{f}\tilde{g}$, it follows that

$$\|\tilde{H} - H\|_1 \leq 2^{-2s} (n+1)^2 \cdot 2^{s+\tau+2} \leq 2^{\tau+2\log(n+1)+2-s} \leq 2^{-\ell},$$

hence an ℓ -bit-approximation as required can be recovered from the exact product of \tilde{f} and \tilde{g} by mere bitshifts. Since $\|\tilde{f}\|_\infty, \|\tilde{g}\|_\infty \leq 2^{s+\tau}$, multiplication of \tilde{f} and \tilde{g} can be carried out exactly in $O(\text{Mul}((s+\tau)n))$ bit operations. This proves the complexity result. For the precision requirement, notice that $\|f\|_\infty, \|g\|_\infty \leq 2^{s+\tau}$; thus, we need $(s+\tau+\lceil \log(n+1) \rceil + 3) \leq (\ell + 2\tau + 3\log n + 8)$ -bit-approximations of the coefficients of F and G to compute \tilde{f} and \tilde{g} . \square

Notice that the norm of the polynomial factors affects the absolute precision of their product. Hence, in order to evaluate the accuracy of the remainders in Algorithm 17, we need good estimates on the norm of Q and R in Algorithm 18. A naive bound for the coefficient growth in step 2 of Algorithm 18 turns out to be too pessimistic. Instead, we give a slightly generalized version of a result due to Schönhage [47, Thm. 4.1]:

Lemma 20. *Let $F = \sum_{i=0}^{2n} F_i x^i \in \mathbb{C}[x]$ be a polynomial of degree at most $2n$ and $G = \sum_{i=0}^n G_i x^i \in \mathbb{C}[x]$ be a monic polynomial of degree n . Let $\rho \geq 1$ be an upper bound on the magnitude of the roots of G . If Q and R are the quotient and the remainder of the division of F by G , that is $F = Q \cdot G + R$ with R uniquely defined by $\deg R < \deg G = n$, then it holds that*

$$\|Q\|_1 \leq 2^{2n} \rho^n \cdot \|F\|_1 \quad \text{and} \quad \|R\|_1 \leq 2^{4n} \rho^{4n} \cdot \|F\|_1.$$

Proof. The coefficients Q_k of Q appear as the leading coefficients in the Laurent series of the function

$$\frac{F(x)/x^n}{G(x)} = \frac{F_{2n} + F_{2n-1}/x + F_{2n-2}/x^2 + \dots}{G_n + G_{n-1}/x + G_{n-2}/x^2 + \dots} = Q_n + \frac{Q_{n-1}}{x} + \frac{Q_{n-2}}{x^2} + \dots,$$

and thus, using Cauchy's integral formula, they can be represented as

$$Q_k = \frac{1}{2\pi i} \int_{|x|=\varrho} \frac{F(x)/x^n}{G(x)} x^{n-k-1} dx \quad (22)$$

with an arbitrary positive number $\varrho > \rho$; see [47, (4.7)–(4.9)]. For an arbitrary x on the boundary of the disk $D_\varrho(0)$ with radius ϱ centered at the origin, it holds that $|F(x) x^{-k-1}| \leq \|F\|_1 \cdot (2\varrho)^{2n-k-1} \leq \|F\|_1 \cdot (2\varrho)^{2n-1}$ and $|G(x)| \geq (\varrho - \rho)^n$ because G is monic and the distance from x to any root of F is at least $\varrho - \rho$. Hence, substitution of $\varrho = 2\rho$ in (22) yields

$$|Q_k| \leq 2^n \cdot (2\rho)^n \cdot \|F\|_1.$$

This proves the bound on $\|Q\|_1$. The second claim now follows from $\|R\|_1 \leq \|F\|_1 + \|Q\|_1 \cdot \|G\|_1$, using the triangle inequality and the submultiplicativity of $\|\cdot\|_1$, and the observation that the magnitude of the i -th coefficient of G is bounded by $\binom{n}{i} \rho^i$ and, hence, $\|G\|_1 \leq \sum \binom{n}{i} \rho^i = (1 + \rho)^n \leq (2\rho)^n$. \square

We are now in the position to assemble the main results for the bit complexity of approximate division with remainder for monic divisors and, in turn, of multipoint evaluation.

Corollary 21. *Let $F \in \mathbb{C}[x]$ be a polynomial of magnitude bounded by $(2n, \tau)$ and $G \in \mathbb{C}[x]$ be a monic polynomial of degree n with a given upper bound $2^\Gamma \geq 1$ on the magnitude of the roots of G , where $\Gamma \geq 1$. Computing an ℓ -bit approximation of $R = F \bmod G$, that is, computing an $\tilde{R} \in \mathbb{C}[x]$ such that $\|\tilde{R} - R\|_1 \leq 2^{-\ell}$, is possible in*

$$\tilde{O}(n(\ell + \tau + n\Gamma))$$

bit operations and with a precision demand of at most $\ell + \tilde{O}(\tau + n\Gamma)$ bits on each of the coefficients of F and G .

Proof. We apply Algorithm 18 in approximate arithmetic with operands $\tilde{\cdot}$ and, if necessary, discard the leading terms of the result \tilde{R} of degree higher than $n - 1$. The analysis of the required precision is done in a backwards fashion. Recall that $\|G\|_1 \leq (1 + 2^\Gamma)^n \leq 2^{O(n\Gamma)}$.

Lemma 19, applied to the computation of R in Step 3, shows that the computation of \tilde{R} entails an error of at most $2^{-\ell}$ if \tilde{F} , \tilde{G} and \tilde{Q} are L -approximations to the exact polynomials for some $L = \ell + \tilde{O}(\tau + n\Gamma)$. Accordingly, in Step 2, we need to obtain an L' -approximation of h_k with some $L' = L + \tilde{O}(\tau + \log \|h_k\|_1)$.

Observe that $\text{rev}_n h_k$ is exactly the quotient of the division of x^{2n} by G , and thus, $\log \|h_k\|_1 = O(n\Gamma)$ by Lemma 20. It follows that $\log \|h_i\|_1 = O(n\Gamma)$ for all $i = 0, \dots, k$ as the intermediate inverses h_i are identical to $h_k \bmod x^{2^i}$. We conclude that each of the k iterations during Step 2 deteriorates the accuracy by at most $O(n\Gamma + \log n)$ bits. Hence, it follows that it suffices to consider L'' -bit approximations of F and G , with $L'' = L' + \lceil \log(n + 1) \rceil \cdot O(n\Gamma + \log n) = \ell + \tilde{O}(\tau + n\Gamma)$, in order to eventually obtain an ℓ -bit approximation for \tilde{R} .

For the bit complexity bound, notice that it suffices to run Algorithm 18 with fixed precision arithmetic and an accuracy of L'' bits after the binary point, where each of the $\tilde{O}(n)$ field operations in \mathbb{C} requires $\tilde{O}(L'') = \tilde{O}(\ell + \tau + n\Gamma)$ bit operations. \square

Theorem 22. *Let $F \in \mathbb{C}[x]$ be a polynomial of magnitude (n, τ) , with $\tau \geq 1$, and let $x_1, \dots, x_n \in \mathbb{C}$ be complex points with absolute values bounded by 2^Γ , where $\Gamma \geq 1$. Then, for an arbitrary positive integer L , the computation of values \tilde{y}_j with $|\tilde{y}_j - F(x_j)| \leq 2^{-L}$ for all j , needs*

$$\tilde{O}(n(L + \tau + n\Gamma))$$

bit operations. Moreover, the precision demand on F and the points x_j is bounded by $L + \tilde{O}(\tau + n\Gamma)$ bits.

Proof. Define $g_{i,j}$ and $r_{i,j}$ as in Algorithm 17. We analyse a run of the algorithm using approximate multiplication and division, with a precision of ℓ_i^{div} for the approximate divisors $\tilde{g}_{i,*}$ and remainders $\tilde{r}_{i,*}$ in the i -th layer of the subproduct and the remainder tree. We recall that $\deg \tilde{g}_{i,*} = \deg g_{i,*} = 2^i$.

According to Corollary 21, for the recursive divisions to yield an output precision $\ell_i \geq 0$, it suffices to have approximations $\tilde{r}_{i+1,*}$ and $\tilde{g}_{i,*}$ of the exact polynomials $f := r_{i+1,*}$ and $g := g_{i,*}$ to a precision of

$$\ell_{i+1}^{\text{div}} := \ell_i^{\text{div}} + \tilde{O}(\log \|r_{i+1,*}\|_1 + 2^i \Gamma) \quad (23)$$

bits, since the roots of each $g_{i,*}$ are contained in the set $\{x_1, \dots, x_n\}$ and, thus, their absolute values are also bounded by 2^Γ . A bound the magnitude of the remainders $r_{i,*}$ for $i < \log n$ is a consequence of Lemma 20, applied in an iterative manner, which yields

$$\log \|r_{i,*}\|_1 = \log \|r_{i+1,*}\|_1 + O(2^{i+1} \Gamma + i \cdot 2^i) = O(\tau + 2n\Gamma + n \log n). \quad (24)$$

For the last estimation, we use that $\|r_{\log n,0}\|_1 = \|F\|_1 \leq (n+1)2^\tau$.

Combining (23) and (24) yields $\ell^{\text{div}} := \max_{i>0} \ell_i^{\text{div}} = \ell_0^{\text{div}} + O(\tau + 2n\Gamma + n \log n)$. Hence, choosing $\ell_0^{\text{div}} := L$, we eventually achieve evaluation up to an error of 2^{-L} if all numerical divisions are carried out with precision ℓ^{div} . The bit complexity to carry out a single numerical division at the i -th layer of the tree is then bounded by $\tilde{O}(2^i(\ell^{\text{div}} + \tau + 2^i \Gamma)) = \tilde{O}(2^i(L + n\Gamma + \tau))$. Since there are $n/2^i$ divisions, the total cost at the i -th layer is bounded by $\tilde{O}(n(L + n\Gamma + \tau))$. The depth of the tree equals $\log n$, and thus the overall bit complexity is $\tilde{O}(n(L + n\Gamma + \tau))$.

It remains to bound the precision demand and, hence, the cost for computing $(L + \tau + 2n\Gamma + O(n \log n))$ -bit approximations of the polynomials $g_{i,*}$. According to Lemma 19, in order to compute the polynomials $g_{i,*}$ to a precision of ℓ_i^{mul} , we have to consider ℓ_{i-1}^{mul} -bit approximations of $g_{i-1,*}$, where

$$\ell_i^{\text{mul}} = \ell_{i-1}^{\text{mul}} + O(\log \|g_{i-1,*}\|_1 + i) = \ell_{i-1}^{\text{mul}} + O(i\Gamma) = \ell_0^{\text{mul}} + \tilde{O}(\log n \cdot \Gamma).$$

Hence, it suffices to run all multiplications in the product tree with a precision of $\ell^{\text{mul}} = L + \tilde{O}(\tau + n\Gamma + n \log n)$. The bit complexity for all multiplications is bounded by $\tilde{O}(n \ell^{\text{mul}}) = \tilde{O}(n(L + \tau + n\Gamma))$, and the precision demand for the points x_i is bounded by $\ell^{\text{mul}} + O(\Gamma + \log n) = L + \tilde{O}(\tau + n\Gamma + n \log n)$. \square

References

- [1] Lionel Alberti, Bernard Mourrain, and Julien Wintz. “Topology and arrangement computation of semi-algebraic planar curves”. In: *Computer Aided Geometric Design* 25.8 (2008), pp. 631–651. ISSN: 0167-8396. DOI: [10.1016/j.cagd.2008.06.009](https://doi.org/10.1016/j.cagd.2008.06.009).
- [2] Dennis S. Arnon and Scott McCallum. “A Polynomial-time Algorithm for the Topological Type of a Real Algebraic Curve”. In: *Journal of Symbolic Computation* 5.1-2 (1988), pp. 213–236. ISSN: 0747-7171. DOI: [10.1016/S0747-7171\(88\)80013-0](https://doi.org/10.1016/S0747-7171(88)80013-0).
- [3] Saugata Basu, Richard Pollack, and Marie-Françoise Roy. *Algorithms in Real Algebraic Geometry*. 2nd. Vol. 10. Algorithms and Computation in Mathematics. Springer, 2006. ISBN: 978-3-540-33098-1.
- [4] Daniel J. Bates et al. *Numerically Solving Polynomial Systems with Bertini*. Software, Environments, and Tools. Society for Industrial and Applied Mathematics, 2013. ISBN: 978-1-611-97269-6.
- [5] Carlos Beltrán and Anton Leykin. “Robust Certified Numerical Homotopy Tracking”. In: *Foundations of Computational Mathematics* 13.2 (2013), pp. 253–295. DOI: [10.1007/s10208-013-9143-2](https://doi.org/10.1007/s10208-013-9143-2).
- [6] Eric Berberich, Pavel Emeliyanenko, and Michael Sagraloff. “An Elimination Method for Solving Bivariate Polynomial Systems: Eliminating the Usual Drawbacks”. In: *Proceedings of the 13th Workshop on Algorithm Engineering and Experiments*. Ed. by Matthias Müller-Hannemann and Renato Fonseca F. Werneck. Society for Industrial and Applied Mathematics, 2011, pp. 35–47. DOI: [10.1137/1.9781611972917.4](https://doi.org/10.1137/1.9781611972917.4).
- [7] Eric Berberich, Michael Hemmer, and Michael Kerber. “A Generic Algebraic Kernel for Non-linear Geometric Applications”. In: *Proceedings of the 27th Symposium on Computational Geometry*. New York, NY, USA: Association for Computing Machinery, 2011, pp. 179–186. DOI: [10.1145/1998196.1998224](https://doi.org/10.1145/1998196.1998224).
- [8] Eric Berberich et al. “Exact symbolic-numeric computation of planar algebraic curves”. In: *Theoretical Computer Science* 491 (2013), pp. 1–32. ISSN: 0304-3975. DOI: [10.1016/j.tcs.2013.04.014](https://doi.org/10.1016/j.tcs.2013.04.014).
- [9] Dario A. Bini and Giuseppe Fiorentino. “Design, analysis, and implementation of a multiprecision polynomial rootfinder”. In: *Numerical Algorithms* 23.2-3 (2000), pp. 127–173. ISSN: 1017-1398. DOI: [10.1023/A:1019199917103](https://doi.org/10.1023/A:1019199917103).
- [10] Y. Bouzidi et al. *Improved algorithm for computing separating linear forms for bivariate systems*. To appear in the Proceedings of the 39th International Symposium on Symbolic and Algebraic Computation. Kobe, Japan, 2014. ARXIV: [1405.4740](https://arxiv.org/abs/1405.4740).
- [11] Yacine Bouzidi et al. “Rational Univariate Representations of Bivariate Systems and Applications”. In: *Proceedings of the 38th International Symposium on Symbolic and Algebraic Computation*. Boston, Maine, USA: Association for Computing Machinery, 2013, pp. 109–116. ISBN: 978-1-4503-2059-7. DOI: [10.1145/2465506.2465519](https://doi.org/10.1145/2465506.2465519).
- [12] Yacine Bouzidi et al. “Separating Linear Forms for Bivariate Systems”. In: *Proceedings of the 38th International Symposium on Symbolic and Algebraic Computation*. Boston, Maine, USA: Association for Computing Machinery, 2013, pp. 117–124. ISBN: 978-1-4503-2059-7. DOI: [10.1145/2465506.2465518](https://doi.org/10.1145/2465506.2465518).

- [13] Michael Burr et al. "Complete subdivision algorithms, II: Isotopic meshing of singular algebraic curves". In: *Journal of Symbolic Computation* 47.2 (2012), pp. 131–152. DOI: [10.1016/j.jsc.2011.08.021](https://doi.org/10.1016/j.jsc.2011.08.021).
- [14] Jin-San Cheng, Xiao-Shan Gao, and Jia Li. "Root Isolation for Bivariate Polynomial Systems with Local Generic Position Method". In: *Proceedings of the 34th International Symposium on Symbolic and Algebraic Computation*. Seoul, Republic of Korea: Association for Computing Machinery, 2009, pp. 103–110. ISBN: 978-1-60558-609-0. DOI: [10.1145/1576702.1576719](https://doi.org/10.1145/1576702.1576719).
- [15] Jin-San Cheng and Kai Jin. "A Generic Position Based Method for Real Root Isolation of Zero-Dimensional Polynomial Systems". In: *Computing Research Repository* (2013). ARXIV: [1312.0462](https://arxiv.org/abs/1312.0462).
- [16] Jinsan Cheng et al. "On the Topology of Real Algebraic Plane Curves". In: *Mathematics in Computer Science* 4.1 (2010), pp. 113–137. ISSN: 1661-8270. DOI: [10.1007/s11786-010-0044-3](https://doi.org/10.1007/s11786-010-0044-3).
- [17] George E. Collins. "Quantifier Elimination for Real Closed Fields by Cylindrical Algebraic Decomposition". In: *Proceedings of the 2nd GI Conference on Automata Theory and Formal Languages*. Vol. 6. Lecture Notes in Computer Science. 1975. Reprinted with corrections in: B. F. Caviness and J. R. Johnson (eds.), *Quantifier Elimination and Cylindrical Algebraic Decomposition*, pp. 85–121. Berlin: Springer. 1998, pp. 134–183. ISBN: 978-3-540-07407-6. DOI: [10.1007/3-540-07407-4_17](https://doi.org/10.1007/3-540-07407-4_17).
- [18] Anindya De et al. "Fast Integer Multiplication Using Modular Arithmetic". In: *Proceedings of the 40th Symposium on the Theory of Computing*. New York, NY, USA: Association for Computing Machinery, 2008, pp. 499–506. DOI: [10.1145/1374376.1374447](https://doi.org/10.1145/1374376.1374447).
- [19] Daouda Niang Diatta, Fabrice Rouillier, and Marie-Françoise Roy. *On the computation of the topology of plane curves*. To appear in the Proceedings of the 39th International Symposium on Symbolic and Algebraic Computation. Kobe, Japan, 2014. halarchivesouvertes: [00935728](https://hal.archives-ouvertes.fr/hal-00935728).
- [20] Dimitrios I. Diochnos, Ioannis Z. Emiris, and Elias P. Tsigaridas. "On the Asymptotic and Practical Complexity of Solving Bivariate Systems over the Reals". In: *Journal of Symbolic Computation* 44.7 (July 2009), pp. 818–835. ISSN: 0747-7171. DOI: [10.1016/j.jsc.2008.04.009](https://doi.org/10.1016/j.jsc.2008.04.009).
- [21] Arno Eigenwillig and Michael Kerber. "Exact and Efficient 2D-Arrangements of Arbitrary Algebraic Curves". In: *Proceedings of the 19th Symposium on Discrete Algorithms*. Ed. by Shang-Hua Teng. San Francisco, California: Association for Computing Machinery & Society for Industrial and Applied Mathematics, 2008, pp. 122–131. ACMID: [1347082.1347096](https://doi.org/10.1145/1347082.1347096).
- [22] Pavel Emeliyanenko. "Computing resultants on Graphics Processing Units: Towards GPU-accelerated computer algebra". In: *Journal of Parallel and Distributed Computing* 73.11 (Nov. 2013), pp. 1494–1505. ISSN: 0743-7315 (print), 1096-0848 (electronic). DOI: [10.1016/j.jpdc.2012.07.015](https://doi.org/10.1016/j.jpdc.2012.07.015).
- [23] Pavel Emeliyanenko. "Harnessing the Power of GPUs for Problems in Real Algebraic Geometry". PhD thesis. Saarbrücken, Germany: University of Saarland, 2012. URN: [urn:nbn:de:bsz:291-scidok-49953](https://nbn-resolving.org/urn:nbn:de:bsz:291-scidok-49953).

- [24] Pavel Emeliyanenko and Michael Sagraloff. “On the Complexity of Solving a Bivariate Polynomial System”. In: *Proceedings of the 37th International Symposium on Symbolic and Algebraic Computation*. Ed. by Joris van der Hoeven and Mark van Hoeij. New York, NY, USA: Association for Computing Machinery, 2012, pp. 154–161. ISBN: 978-1-4503-1269-1. DOI: [10.1145/2442829.2442854](https://doi.org/10.1145/2442829.2442854).
- [25] Efi Fogel, Dan Halperin, and Ron Wein. *CGAL Arrangements and Their Applications: A Step-by-Step Guide*. Vol. 7. Geometry and Computing. Springer, 2012. ISBN: 978-3-642-17282-3. DOI: [10.1007/978-3-642-17283-0](https://doi.org/10.1007/978-3-642-17283-0).
- [26] Martin Fürer. “Faster Integer Multiplication”. In: *Journal on Computing* 39.3 (2009), pp. 979–1005. DOI: [10.1137/070711761](https://doi.org/10.1137/070711761).
- [27] Joachim von zur Gathen and Jürgen Gerhard. *Modern Computer Algebra*. 3rd. Cambridge, UK: Cambridge University Press, 2013. ISBN: 978-1-107-03903-2.
- [28] Izrail Moiseevitch Gelfand, Mikhail M. Kapranov, and Andrei V. Zelevinsky. *Discriminants, resultants, and multidimensional determinants*. Mathematics: Theory and Application. Basel: Birkhäuser, 1994. ISBN: 0-8176-3660-9. DOI: [10.1007/978-0-8176-4771-1](https://doi.org/10.1007/978-0-8176-4771-1).
- [29] Laureano González-Vega and M’Hammed El Kahoui. “An Improved Upper Complexity Bound for the Topology Computation of a Real Algebraic Plane Curve”. In: *Journal of Complexity* 12.4 (Dec. 1996), pp. 527–544. ISSN: 0885-064X. DOI: [10.1006/jcom.1996.0032](https://doi.org/10.1006/jcom.1996.0032).
- [30] Janusz Gwoździewicz and Arkadiusz Płoski. “Formulae for the Singularities at Infinity of Plane Algebraic Curves”. In: *Acta Mathematica* 1255(39) (2001), pp. 109–133. URL: <http://eudml.org/doc/122736>.
- [31] Michael Kerber. “Geometric Algorithms for Algebraic Curves and Surfaces”. PhD thesis. Saarbrücken, Germany: University of Saarland, 2009. URN: [urn:nbn:de:bsz:291-scidok-29490](https://nbn-resolving.org/urn:nbn:de:bsz:291-scidok-29490).
- [32] Michael Kerber and Michael Sagraloff. “A worst-case bound for topology computation of algebraic curves”. In: *Journal of Symbolic Computation* 47.3 (2012), pp. 239–258. DOI: [10.1016/j.jsc.2011.11.001](https://doi.org/10.1016/j.jsc.2011.11.001).
- [33] Peter Kirrinnis. “Partial Fraction Decomposition in $\mathbb{C}(z)$ and Simultaneous Newton Iteration for Factorization in $\mathbb{C}[z]$ ”. In: *Journal of Complexity* 14.3 (1998), pp. 378–444. DOI: [10.1006/jcom.1998.0481](https://doi.org/10.1006/jcom.1998.0481).
- [34] Alexander Kobel. “Certified Numerical Root Finding”. MA thesis. Saarbrücken, Germany: University of Saarland, 2011.
- [35] Alexander Kobel and Michael Sagraloff. “Fast Approximate Polynomial Multipoint Evaluation and Applications”. In: *Computing Research Repository* (2013). ARXIV: [1304.8069](https://arxiv.org/abs/1304.8069).
- [36] Tien-Yien Li. “Numerical Solution of Polynomial Systems by Homotopy Continuation Methods”. In: *Handbook of Numerical Analysis*. Vol. 11. Elsevier, 2003, pp. 209–304. DOI: [10.1016/S1570-8659\(02\)11004-0](https://doi.org/10.1016/S1570-8659(02)11004-0).

- [37] Thomas Lickteig and Marie-Françoise Roy. “Sylvester-Habicht Sequences and Fast Cauchy Index Computation”. In: *Journal of Symbolic Computation* 31.3 (2001), pp. 315–341. ISSN: 0747-7171. DOI: [10.1006/jsc.2000.0427](https://doi.org/10.1006/jsc.2000.0427).
- [38] O. P. Lossers. “A Hadamard-Type Bound on the Coefficients of a Determinant of Polynomials”. In: *SIAM Review* 16.3 (1974). Solution to an exercise by A. Jay Goldstein and Ronald L. Graham, pp. 394–395. DOI: [10.1137/1016065](https://doi.org/10.1137/1016065).
- [39] Kurt Mehlhorn, Michael Sagraloff, and Pengming Wang. “From Approximate Factorization to Root Isolation with Application to Cylindrical Algebraic Decomposition”. In: *Proceedings of the 38th International Symposium on Symbolic and Algebraic Computation*. Full version available in the *Computing Research Repository*. Ed. by Manuel Kauers. Boston, Maine, USA: Association for Computing Machinery, 2013, pp. 283–290. ISBN: 978-1-4503-2059-7. DOI: [10.1145/2465506.2465523](https://doi.org/10.1145/2465506.2465523).
- [40] Robert T. Moenck and Allan Borodin. “Fast modular transforms via division”. In: *Proceedings of the 13th Symposium on Switching and Automata Theory*. IEEE Computer Society, 1972, pp. 90–96. DOI: [10.1109/SWAT.1972.5](https://doi.org/10.1109/SWAT.1972.5).
- [41] Bernard Mourrain and Jean-Pascal Pavone. “Subdivision Methods for Solving Polynomial Equations”. In: *Journal of Symbolic Computation* 44.3 (Mar. 2009), pp. 292–306. ISSN: 0747-7171. DOI: [10.1016/j.jsc.2008.04.016](https://doi.org/10.1016/j.jsc.2008.04.016).
- [42] Victor Y. Pan. “Univariate Polynomials: Nearly Optimal Algorithms for Numerical Factorization and Root-finding”. In: *Journal of Symbolic Computation* 33.5 (2002), pp. 701–733. ISSN: 0747-7171. DOI: [10.1006/jsc.2002.0531](https://doi.org/10.1006/jsc.2002.0531).
- [43] Simon Plantinga and Gert Vegter. “Isotopic Approximation of Implicit Curves and Surfaces”. In: *Proceedings of the 2nd Symposium on Geometry Processing*. Nice, France: Association for Computing Machinery, 2004, pp. 245–254. ISBN: 3-905673-13-4. DOI: [10.1145/1057432.1057465](https://doi.org/10.1145/1057432.1057465).
- [44] Daniel Reischert. “Asymptotically Fast Computation of Subresultants”. In: *Proceedings of the 10th International Symposium on Symbolic and Algebraic Computation*. Maui, Hawaii, USA: Association for Computing Machinery, 1997, pp. 233–240. ISBN: 0-89791-875-4. DOI: [10.1145/258726.258792](https://doi.org/10.1145/258726.258792).
- [45] Fabrice Rouillier. “On Solving Systems of Bivariate Polynomials”. In: *Proceedings of the 3rd International Congress on Mathematical Software*. Ed. by Komei Fukuda et al. Lecture Notes in Computer Science. Springer, 2010, pp. 100–104. ISBN: 978-3-642-15581-9. DOI: [10.1007/978-3-642-15582-6_21](https://doi.org/10.1007/978-3-642-15582-6_21).
- [46] Michael Sagraloff and Kurt Mehlhorn. “Computing Real Roots of Real Polynomials: An Efficient Method Based on Descartes’ Rule of Signs and Newton Iteration”. In: *Computing Research Repository* (2013). ARXIV: [1308.4088](https://arxiv.org/abs/1308.4088).
- [47] Arnold Schönhage. “Asymptotically Fast Algorithms for the Numerical Multiplication and Division of Polynomials with Complex Coefficients”. In: *Computer Algebra*. Ed. by Jacques Calmet. Vol. 144. Lecture Notes in Computer Science. Springer Berlin Heidelberg, 1982, pp. 3–15. ISBN: 978-3-540-11607-3. DOI: [10.1007/3-540-11607-9_1](https://doi.org/10.1007/3-540-11607-9_1).
- [48] Arnold Schönhage and Volker Strassen. “Schnelle Multiplikation großer Zahlen”. German. In: *Computing* 7.3-4 (1971), pp. 281–292. ISSN: 0010-485X. DOI: [10.1007/BF02242355](https://doi.org/10.1007/BF02242355).

- [49] Andrew J. Sommese and Charles W. Wampler. *The Numerical Solution of Systems of Polynomials Arising in Engineering and Science*. World Scientific, 2005. ISBN: 978-9-812-56772-7.
- [50] Adam W. Strzeboński. “Cylindrical Algebraic Decomposition Using Validated Numerics”. In: *Journal of Symbolic Computation* 41.9 (2006), pp. 1021–1038. ISSN: 0747-7171. DOI: [10.1016/j.jsc.2006.06.004](https://doi.org/10.1016/j.jsc.2006.06.004).
- [51] Bernard Teissier. “Cycles évanescents, sections planes et conditions de Whitney”. French. In: *Singularités à Cargèse*. Astérisque 7-8 (1973), pp. 285–362.
- [52] Jan Verschelde. “Polynomial Homotopy Continuation with PHCpack”. In: *Communications in Computer Algebra* 44 (3/4 Jan. 2011), pp. 217–220. ISSN: 1932-2240. DOI: [10.1145/1940475.1940524](https://doi.org/10.1145/1940475.1940524).
- [53] Chee Keng Yap. *Fundamental Problems of Algorithmic Algebra*. New York, NY, USA: Oxford University Press, 2000. ISBN: 0-19-512516-9.